

Mathematical Modeling of Competitive Group Recommendation Systems with Application to Peer Review Systems

Hong Xie John C.S. Lui
 Computer Science & Engineering Department
 The Chinese University of Hong Kong
 Email: {hxie,cslui}@cse.cuhk.edu.hk

March 2, 2013

Abstract

In this paper, we present a mathematical model to capture various factors which may influence the accuracy of a competitive group recommendation system. We apply this model to peer review systems, i.e., conference or research grants review, which is an essential component in our scientific community. We explore number of important questions, i.e., how will the number of reviews per paper affect the accuracy of the overall recommendation? Will the score aggregation policy influence the final recommendation? How reviewers' preference may affect the accuracy of the final recommendation? To answer these important questions, we formally analyze our model. Through this analysis, we obtain the insight on how to design a randomized algorithm which is both computationally efficient and asymptotically accurate in evaluating the accuracy of a competitive group recommendation system. We obtain number of interesting observations: i.e., for a medium tier conference, three reviews per paper is sufficient for a high accuracy recommendation. For prestigious conferences, one may need at least seven reviews per paper to achieve high accuracy. We also propose a heterogeneous review strategy which requires equal or less reviewing workload, but can improve over a homogeneous review strategy in recommendation accuracy by as much as 30%. We believe our models and methodology are important building blocks to study competitive group recommendation systems.

1 Introduction

In recent years, recommendation systems [20] have received a lot of attention in both commercial and academic communities. Researchers investigate various algorithmic and complexity issues [7, 10, 20, 22, 23], at the same time, we also see successful applications of recommendation systems in commercial products. In general, recommendation systems take into account a user's preference and make a recommendation so as to maximize the user's utility. *Group recommendation systems* [3], on the other hand, take into account the preferences of *all* users in a group to make a single recommendation. In recent years, we have seen successful group recommendations in commercial areas [1, 3, 12, 13, 15, 16, 18].

In this paper, we consider a special class of recommendation system which we call the *competitive group recommendation system*: there are N users and the system will make a single recommendation to k users only, where $k \leq N$, while $N - k$ users will receive the *complement* of the recommendation. Competitive group recommendation systems have many important applications. In here, we consider an application which is dearest to many researchers' heart: *peer review systems* for conferences or research grant proposals. To a certain degree, the progress of our scientific community depends on the accuracy this type of recommendation systems. To the best of our knowledge, this is the first paper which provides a formal mathematical analysis to such recommendation systems.

A peer review system can be briefly described as follows: there are N candidates (papers or grant proposals), a group of reviewers is asked to review these candidates. Each reviewer evaluates a subset of these candidates based on her preference, and will provide a rating for each candidate. The system will use some policies to aggregate all ratings of all candidates, and will only recommend a subset k candidates for acceptance, while all other candidates will receive a rejection, which is the complement of the acceptance recommendation. For such systems, there are many interesting

questions to explore, e.g., to achieve high accuracy, how many reviews each candidate should receive? What is the probability that the best candidate will be accepted or rejected? How reviewers' preference may influence the final recommendation? Is one rating aggregation policy more accurate than others?

Our contribution can be summarized as follows:

- We propose a mathematical model to understand the accuracy of a competitive group recommendation system and apply it to conference review systems.
- We formally analyze the model. Through this analysis, we gain the insight to create a randomized algorithm to evaluate the model. We show our algorithm is computationally efficient and also provides performance guarantees.
- We apply our model to a conference review system and show many interesting insights, i.e., for a medium titer conference, three reviews per paper can guarantee a highly accurate recommendation, but for prestigious conferences, we need at least seven reviews per paper.
- We propose a two round heterogeneous review strategy which outperforms the homogeneous review strategy by as much as 30% in recommendation accuracy with the the same or less reviewing workload.

This is the outline of the paper. In Section 2, we present the mathematical model of competitive group recommendation systems. In Section 3, we present analysis and derive theoretical results of the model. In Section 4, we propose a randomized algorithm which is computationally efficient and provides performance guarantees in evaluating a competitive group recommendation system. In Section 5, we evaluate the performance of a conference review system and explore various factors that influence its accuracy. Related work is given in Section 6 and Section 7 concludes.

2 Mathematical Model

Let us present the mathematical model of a competitive group recommendation system and we focus on a particular application scenario, a conference review system which is a representative example of peer review systems. Let $\mathcal{P} = \{P_1, \dots, P_N\}$ be a finite set of N candidate papers. Let $Q_i \in (1, m)$ represent the *intrinsic quality* of paper P_i . Higher value of intrinsic quality implies higher quality. Hence, if $Q_i > Q_j$, it means paper P_i is better than P_j . Without any loss of generality, let us assume $Q_1 > Q_2 > \dots > Q_N$. It is important to emphasize that reviewers of these papers do *not* have any a-prior knowledge of $Q_i, \forall i$. The conference can only accept k papers, where $1 \leq k \leq N$. Let $\mathcal{A}^I(k)$ and $\mathcal{A}(k)$ denote the set of the k accepted papers according to the intrinsic quality or according to the conference recommendation criteria respectively. It is clear that $\mathcal{A}^I(k) = \{P_1, P_2, \dots, P_k\}$, and if a conference recommendation system is perfect, we should have $\mathcal{A}^I(k) = \mathcal{A}(k)$. But in general, many factors or reviewers' preference may influence the final recommendation, hence $\mathcal{A}^I(k) \neq \mathcal{A}(k)$. To measure the accuracy of a recommendation system, we aim to determine how many papers in $\mathcal{A}(k)$ are also in $\mathcal{A}^I(k)$. Formally, we seek to derive the following probability mass function (pmf) :

$$\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = i], \quad \text{for } i = 0, 1, \dots, k.$$

Intuitively, if $\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = k]$ occurs with a high probability, then the conference recommendation system is *accurate* and at the same time, robust against different scoring and human factors.

Let \mathcal{R} be a finite set of M reviewers. We assume that the reviewers are independent. Reviewers do not have a direct knowledge of $Q_i, \forall i$, the intrinsic quality of papers, and they evaluate papers based on their own preference. A reviewer submits a score for each paper after reviewing. Scores are discrete and take on value in $\{1, \dots, m\}$. Paper P_i , for $i = 1, \dots, N$, is assigned to $n_i \geq 1$ reviewers. Hence, paper P_i receives n_i reviewing scores. Let $\mathcal{S}^i = \{S_1^i, \dots, S_{n_i}^i\}$ denote the set of n_i scores of paper P_i , where $S_j^i \in \{1, \dots, m\}$. Let $\mathcal{R}(S_j^i)$ represent the reviewer who submits score S_j^i . Let $e_j^i \in \{1, \dots, l\}$ represent the expertise level (or familiarity) that reviewer $\mathcal{R}(S_j^i)$ selects on topics related to paper P_i . Reviewer $\mathcal{R}(S_j^i)$ submits score S_j^i in conjunction with expertise level e_j^i . Again, we adopt the convention that higher values represent higher quality or expertise level. There are number of interesting questions one can explore, i.e., how M , the number of reviewers (or the size of a technical program committee), as well as $n_i, \forall i$, the number of reviews for each paper, may affect the accuracy of the final recommendation?

Let \mathcal{V} be the voting rule that is used by the conference recommendation system to rank papers based on their reviewing scores. Generally, a voting rule works in two steps. It first aggregates the reviewing scores of each paper into

a combined overall score. Then it ranks all papers based on their respective combined overall scores. Let $\gamma_i = \mathcal{V}(\mathcal{S}^i)$ be the combined overall score of paper P_i derived from \mathcal{S}^i under the voting rule \mathcal{V} . There can be many voting rules. A simple and often used voting rule is the *average score rule*. In this case, we have $\gamma_i = \sum_{S \in \mathcal{S}^i} S/|\mathcal{S}^i|$. By sorting $\gamma_1, \dots, \gamma_N$, we obtain a ranked list of all papers. Again, there are number of interesting questions to explore, i.e., what are some effective voting rules? Can one voting rule be more accurate than others?

Specifying the voting rule is not enough. Recall that the system can only accept k papers. It may happen that the combined overall score of the k^{th} ranked paper, equals to that of $(k+1)$ -th ranked paper. In this case, we need to specify a *tie-breaking rule* to decide which paper should be selected. Let \mathcal{T} denote the tie-breaking rule. It is interesting to explore whether the recommendation results are sensitive to a particular tie-breaking rule.

To answer the above questions, let us now present probabilistic models in describing the intrinsic quality (or the *self-selection effect*), the reviewing behavior, as well as critical degree of reviewers.

2.1 Model Intrinsic Quality via Self-selection

It is well-known that paper submission has the *self-selection effect*. In other words, authors tend to submit their high quality papers to some highly prestigious and selective conferences, while lower tier conferences may receive papers with lower quality, or candidate papers have high variance in quality. The intrinsic quality of a submitted paper can be described as a random variable, and one can vary its mean or variance to reflect the *self-selection effect*. Specifically, a high value of mean and a small value of variance imply that the submitted papers are of high intrinsic qualities and these qualities have small variation only. On the other hand, a low value of mean and a large value of variance imply that the submitted papers have low intrinsic qualities and these qualities have high variability.

We use $Q_i \in (1, m)$ to denote the intrinsic quality of paper P_i . Assume Q_1, \dots, Q_N are independent random variables. Let $\mathcal{D}(Q_i)$ denote the probability distribution of Q_i . The probability distribution $\mathcal{D}(Q_i)$ is described by a truncated normal distribution $\mathcal{N}(q_i, \sigma_i^2)$, where $q_i \in (1, m)$ is the mean and σ_i^2 is the variance. Since the value of the intrinsic quality Q_i is in $(1, m)$, thus $\mathcal{D}(Q_i)$ is obtained by truncating $\mathcal{N}(q_i, \sigma_i^2)$ to keep those values in $(1, m)$ and scaling up the kept values by $1/\Pr[1 < X < m]$, where X is a random variable with probability distribution $\mathcal{N}(q_i, \sigma_i^2)$. It should be clear that after truncation, the mean q_i and the variance σ_i^2 can still reflect the *self-selection effect*. Here, we use the following parameters to reflect four representative types of *self-selectivity*:

High self-selectivity: the mean q_i and variance σ_i^2 are specified by

$$q_i = m, \sigma_i^2 = 1, \quad \text{for } i = 1, \dots, N. \quad (1)$$

This indicates that papers tend to have high intrinsic quality (or high mean), and most of the probability mass concentrates around high intrinsic quality. Top tier conferences fall into this category.

Medium self-selectivity: the mean q_i and variance σ_i^2 are

$$q_i = (m+1)/2, \sigma_i^2 = 1, \quad \text{for } i = 1, \dots, N. \quad (2)$$

This reflects that papers tend to have an average intrinsic quality and most of the probability mass concentrates around the average intrinsic quality. Medium tier conferences fall into this category.

Low self-selectivity: the mean q_i and variance σ_i^2 are

$$q_i = 1, \sigma_i^2 = 1, \quad \text{for } i = 1, \dots, N. \quad (3)$$

This indicates that papers tend to have low intrinsic quality (or low mean), and most of the probability mass concentrates around low quality. Low tier conferences fall into this category.

Random self-selectivity: the variance σ_i^2 is

$$\sigma_i^2 = \infty, \quad \text{for } i = 1, \dots, N. \quad (4)$$

So $\mathcal{D}(Q_i)$ converges to a uniform distribution on $(1, m)$. This means that the intrinsic qualities of submitted papers are uniformly distributed. Newly started conferences fall into this category. This is because a newly started conference has not built up a reputation yet, thus researchers are not sure if it is a good conference, which results in random quality in submission.

2.2 Model for Reviewing Behavior

When reviewing a paper, a reviewer needs to evaluate its quality. Here we assume that each reviewer is fair, unbiased and critical. We consider two most important factors that affect the evaluation. The first one is Q_i , the intrinsic quality of paper P_i , and the second one is the critical degree of the reviewer. Specifically, when Q_i is high, the evaluated quality of P_i is more likely to be high. And the higher the critical degree of the reviewer, the more likely that the evaluated quality tends to be close to the intrinsic quality of that paper. The reviewing behavior can be described by a random variable and one can vary its mean and variance to reflect the intrinsic quality and critical degree.

To illustrate, consider a paper P_i and one of its score S_j^i . Recall that the reviewer who submits score S_j^i is denoted by $\mathcal{R}(S_j^i)$. Let $c_j^i \in [0, 1]$ denote the critical degree of reviewer $\mathcal{R}(S_j^i)$. With the usual convention, higher value represents higher critical degree. The score S_j^i is a random variable with probability distribution $\mathcal{D}(S_j^i)$, which should have the following two properties:

Property 1: The mean should be equal to Q_i . The physical meaning is that a reviewer is unbiased.

Property 2: The variance should reflect the critical degree of a reviewer. Specifically, the higher the critical degree, the lower the variance for the probability distribution $\mathcal{D}(S_j^i)$.

In our study, the probability distribution $\mathcal{D}(S_j^i)$ is obtained by mapping a normal distribution $\mathcal{N}(Q_i, \sigma^2(c_j^i))$ to a discrete distribution. Note that the standard variance $\sigma(c_j^i)$ is a monotonic decreasing function of c_j^i and we will specify it in later section. The probability distribution mapping can be described by the following two steps:

Discretization : Transform a normal distribution into a discrete distribution. We transform the normal distribution $\mathcal{N}(Q_i, \sigma^2(c_j^i))$ into a discrete random variable L with probability distribution $\tilde{\mathcal{N}}(Q_i, \sigma^2(c_j^i))$ with values in $\{1, \dots, m\}$. The pmf of L is:

$$\Pr[L = \ell] = \frac{\Pr[\ell - 0.5 \leq X \leq \ell + 0.5]}{\Pr[0.5 \leq X \leq m + 0.5]} = \frac{\Phi((\ell + 0.5 - Q_i)/\sigma(c_j^i)) - \Phi((\ell - 0.5 - Q_i)/\sigma(c_j^i))}{\Phi((m + 0.5 - Q_i)/\sigma(c_j^i)) - \Phi((0.5 - Q_i)/\sigma(c_j^i))},$$

for $\ell = 1, \dots, m$ (5)

where $\Phi(x) = \int_{-\infty}^x \exp(-t^2/2)/\sqrt{2\pi} dt$ and the probability distribution of X is $\mathcal{N}(Q_i, \sigma^2(c_j^i))$. Note that this discrete distribution satisfies Property 2 but not Property 1. In the following step, we adjust the distribution so that it satisfies Property 1 also.

Adjustment: Adjust the distribution $\tilde{\mathcal{N}}(Q_i, \sigma^2(c_j^i))$ such that its mean equals to Q_i . The idea is that if $E[L] < Q_i$, then we increase the mean of $\tilde{\mathcal{N}}(Q_i, \sigma^2(c_j^i))$ by scaling up the probability:

$$\Pr[L = \ell], \quad \text{for all } \ell = \lfloor Q_i \rfloor + 1, \dots, m.$$

Else, we decrease the mean by scaling them down. Applying this idea to adjust the mean of $\tilde{\mathcal{N}}(Q_i, \sigma^2(c_j^i))$, we obtain the probability distribution $\mathcal{D}(S_j^i)$. The pmf of $\mathcal{D}(S_j^i)$ is:

$$\Pr[S_j^i = \ell] = \begin{cases} [1 - \beta(Q_i, c_j^i)]\Pr[L = \ell], & \forall \ell = 1, \dots, \lfloor Q_i \rfloor \\ [1 + \alpha(Q_i, c_j^i)]\Pr[L = \ell], & \forall \ell = \lfloor Q_i \rfloor + 1, \dots, m \end{cases}, \quad (6)$$

where $\alpha(Q_i, c_j^i)$ and $\beta(Q_i, c_j^i)$ are:

$$\alpha(Q_i, c_j^i) = \frac{\sum_{\ell=1}^{\lfloor Q_i \rfloor} \Pr[L = \ell](Q_i - E[L])}{\sum_{\ell=1}^{\lfloor Q_i \rfloor} \Pr[L = \ell](E[L] - \ell)}, \quad (7)$$

$$\beta(Q_i, c_j^i) = \frac{\sum_{\ell=\lfloor Q_i \rfloor+1}^m \Pr[L = \ell](Q_i - E[L])}{\sum_{\ell=1}^{\lfloor Q_i \rfloor} \Pr[L = \ell](E[L] - \ell)}, \quad (8)$$

and L is a discrete random variable with probability distribution $\tilde{\mathcal{N}}(Q_i, \sigma^2(c_j^i))$. Note that $\mathcal{D}(S_j^i)$ satisfies both Property 1 and 2.

2.3 Model for Critical Degree

To model the critical degree of a reviewer, we classify papers and reviewers into “types”. Specifically, a paper can be of many types (e.g., system paper, theory paper, etc), and reviewers can be of many types also (e.g., prefer system paper, or theory paper, etc). If a paper-reviewer pairing is of the same type, then the expertise level and the critical degree of the reviewer will be of high values, else they will be of low values.

To illustrate, consider a paper $P \in \mathcal{P}$ and a reviewer $R \in \mathcal{R}$. Assume reviewer R reviews paper P . Let $u \in [0, 1]$ denote the matching degree between reviewer R and paper P and let e, c denote the corresponding expertise level and critical degree respectively. Using our usual convention, higher value represents higher matching degree. The matching degree couples the expertise level and the critical degree in the following manner:

$$\begin{aligned} e &= \kappa, & \text{if } \mu \in [(\kappa - 1)/l, \kappa/l), \text{ for } \kappa = 1, \dots, l, \\ c &= f(\mu) \in [0, 1], & \text{where } \mu \in [0, 1]. \end{aligned} \quad (9)$$

Note that $e = l$ when $\mu = 1$ and $f(\mu)$ is a monotonic increasing function of μ . There are number of choices for function $f(\mu)$, e.g., $f(\mu) = \mu$, or $f(\mu) = \mu^2$, etc. We will specify it in later section.

Again there are number of interesting questions to explore, i.e., is the conference recommendation system sensitive to the paper-reviewer matching? Will a small percentage of reviewers who prefer theory create a large inaccuracy in the final recommendation of a system-oriented conference (or vice versa)?

3 Theoretical Analysis

Recall that $\mathcal{A}^I(k)$ and $\mathcal{A}(k)$ denote the set of k accepted papers according to the intrinsic quality or according to the conference recommendation criteria respectively. In this section, we first derive the following probability mass function (pmf) :

$$\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = i], \quad \text{for } i = 0, 1, \dots, k.$$

With this pmf, we can then derive the expectation $E[|\mathcal{A}^I(k) \cap \mathcal{A}(k)|]$ and variance $\text{Var}[|\mathcal{A}^I(k) \cap \mathcal{A}(k)|]$. The above probability measures can provide us with a lot of insights, e.g., if $\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = k]$ occurs with a high probability, or $E[|\mathcal{A}^I(k) \cap \mathcal{A}(k)|] \approx k$, then the conference recommendation system is very accurate and robust against different human factors, or if $\text{Var}[|\mathcal{A}^I(k) \cap \mathcal{A}(k)|]$ is of small value, then the conference recommendation system is very stable likely to be close to the expectation. To derive this pmf, let us first consider the following special case. The purpose is to show the general idea of derivation and to illustrate the underlying computational complexity. We will consider the derivation of the general case later.

3.1 Derivation of the Special Case

Let us consider a conference recommendation system which has only one type of papers and one type of reviewers (e.g., all papers are theoretical and all reviewers prefer theoretical topics). Hence, the critical degree of all reviewers are the same, say c . The intrinsic quality of each paper is specified as follows:

$$Q_i = m - i(m - 1)/(N + 1), \quad \text{for } i = 1, \dots, N. \quad (10)$$

Each paper will have the same number of reviews, or $n_i = n, \forall i$. The voting rule \mathcal{V} is the *average score rule* and we use a random rule to tie-break any papers whose scores are the same.

3.1.1 Theoretical derivation

The score set for paper P_i is $\{S_1^i, \dots, S_n^i\}$. Recall that a score is described by a random variable, and its probability distribution is uniquely determined by the intrinsic quality of the corresponding paper and critical degree of the corresponding reviewer. Since the critical degree of each reviewer is the same c , thus S_1^i, \dots, S_n^i are i.i.d. random variables. By specifying Q_i and c_j^i in Eq. (6) with Eq. (10) and $c_j^i = c$ respectively, we obtain the pmf of score S_j^i . This is stated in the following lemma.

Lemma 1 The pmf of score S_j^i , for $i=1, \dots, N$, $j=1, \dots, n$, is

$$Pr[S_j^i = \ell] = \begin{cases} [1 - \beta(m - \frac{i(m-1)}{N+1}, c)]Pr[L = \ell], & \forall \ell = 1, \dots, \lfloor Q_i \rfloor \\ [1 + \alpha(m - \frac{i(m-1)}{N+1}, c)]Pr[L = \ell], & \forall \ell = \lfloor Q_i \rfloor + 1, \dots, m, \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

where L is a discrete random variable with probability distribution $\tilde{N}(m - i(m-1)/(N+1), \sigma^2(c))$ whose pmf is derived by Eq. (5), and $\alpha(m - i(m-1)/(N+1), c)$, $\beta(m - i(m-1)/(N+1), c)$ are derived by Eq. (7) and (8) respectively.

The average score of each paper is $\gamma_i = \sum_{j=1}^n S_j^i/n, \forall i$. The probability mass function (pmf) of the average score of each paper is specified in the following lemma.

Lemma 2 The pmf of the averages score $\gamma_i, \forall i$, is

$$Pr\left[\gamma_i = \frac{\ell}{n}\right] = \sum_{\sum_{j=1}^n s_j = \ell} \prod_{j=1}^n Pr[S_j^i = s_j], \quad \text{for all } \ell = n, \dots, nm, \quad (12)$$

and its cumulative distribution function (CDF) is

$$Pr\left[\gamma_i \leq \frac{\ell}{n}\right] = \sum_{\sum_{j=1}^n s_j \leq \ell} \prod_{j=1}^n Pr[S_j^i = s_j], \quad \text{for all } \ell = n, \dots, nm, \quad (13)$$

where $Pr[S_j^i = s_j]$ is specified in Eq. (11).

Proof: Note that, the ratings of each paper are independent random variables. The distribution of each rating has been derived in Lemma 1. Since $\gamma_i = \sum_{j=1}^n S_j^i/n$, thus by enumerating all the cases satisfying the condition $\sum_{j=1}^n S_j^i = \ell$, we could obtain the pmf of γ_i , or

$$Pr\left[\gamma_i = \frac{\ell}{n}\right] = Pr\left[\sum_{j=1}^n S_j^i = \ell\right] = \sum_{\sum_{j=1}^n s_j = \ell} \prod_{j=1}^n Pr[S_j^i = s_j].$$

Similarly, by enumerating all the cases satisfying the condition $\sum_{j=1}^n S_j^i \leq \ell$, we could obtain the CDF of γ_i , or

$$Pr\left[\gamma_i \leq \frac{\ell}{n}\right] = Pr\left[\sum_{j=1}^n S_j^i \leq \ell\right] = \sum_{\sum_{j=1}^n s_j \leq \ell} \prod_{j=1}^n Pr[S_j^i = s_j],$$

which completes the proof. ■

Based on the results derived in Lemma 1 and Lemma 2, the probability that a specific set of papers is accepted is stated as follows.

Theorem 1 Let $\{P_{i_1}, \dots, P_{i_k}\}$ be a set of k papers. The probability that the accepted paper set equals to $\{P_{i_1}, \dots, P_{i_k}\}$ is:

$$\begin{aligned} Pr[\mathcal{A}(k) = \{P_{i_1}, \dots, P_{i_k}\}] &= \sum_{\ell=n}^{nm} \left(\prod_{i \in \mathcal{I}} Pr[\gamma_i \leq \ell/n] - \prod_{i \in \mathcal{I}} Pr[\gamma_i \leq (\ell-1)/n] \right) \prod_{j \in \bar{\mathcal{I}}} Pr[\gamma_j \leq (\ell-1)/n] + \\ &\quad \sum_{\substack{\mathcal{F} \subseteq \mathcal{I}, \mathcal{G} \subseteq \bar{\mathcal{I}}, \mathcal{F}, \mathcal{G} \neq \emptyset}} \left(\frac{|\mathcal{F} \cup \mathcal{G}|}{|\mathcal{F}|} \right)^{-1} \sum_{\ell=n}^{nm} \prod_{i \in \mathcal{I} \setminus \mathcal{F}} (1 - Pr[\gamma_i \leq \ell/n]) \prod_{j \in \mathcal{F} \cup \mathcal{G}} Pr[\gamma_j = \ell/n] \\ &\quad \prod_{\kappa \in \bar{\mathcal{I}} \setminus \mathcal{G}} Pr[\gamma_\kappa \leq (\ell-1)/n], \end{aligned} \quad (14)$$

where $\mathcal{I} = \{i_1, \dots, i_k\}$ is the index set of $\{P_{i_1}, \dots, P_{i_k}\}$, and $\bar{\mathcal{I}} = \{1, \dots, N\} \setminus \mathcal{I}$ is the complement of \mathcal{I} . And $Pr[\gamma_i = \ell/n]$ is specified in Eq. (12), and $Pr[\gamma_i \leq \ell/n]$ is specified in Eq. (13).

Proof: Let $\mathcal{H} = \{P_{i_1}, \dots, P_{i_k}\}$. Let $\overline{\mathcal{H}} = \{P_1, \dots, P_N\} \setminus \mathcal{H}$ be the complement of \mathcal{H} . Let $\gamma_{\min}(\mathcal{H}) = \min\{\gamma_i, i \in \mathcal{I}\}$ denote the minimum average score of paper set \mathcal{H} . Let $\gamma_{\max}(\overline{\mathcal{H}}) = \max\{\gamma_i, i \in \overline{\mathcal{I}}\}$ denote the maximum average score of paper set $\overline{\mathcal{H}}$. The probability that the accepted paper set equals to $\{P_{i_1}, \dots, P_{i_k}\}$ can be divided into the following three parts:

$$\begin{aligned} \Pr[\mathcal{A}(k) = \{P_{i_1}, \dots, P_{i_k}\}] &= \Pr[\mathcal{A}(k) = \mathcal{H}, \gamma_{\min}(\mathcal{H}) < \gamma_{\max}(\overline{\mathcal{H}})] + \\ &\quad \Pr[\mathcal{A}(k) = \mathcal{H}, \gamma_{\min}(\mathcal{H}) > \gamma_{\max}(\overline{\mathcal{H}})] + \\ &\quad \Pr[\mathcal{A}(k) = \mathcal{H}, \gamma_{\min}(\mathcal{H}) = \gamma_{\max}(\overline{\mathcal{H}})]. \end{aligned} \quad (15)$$

Let us derive these three terms one by one.

According to our voting rule, the accepted paper set $\mathcal{A}(k)$ equals to \mathcal{H} is 0 conditioned on that $\gamma_{\min}(\mathcal{H})$ is less than $\gamma_{\max}(\overline{\mathcal{H}})$. Thus,

$$\begin{aligned} \Pr[\mathcal{A}(k) = \mathcal{H}, \gamma_{\min}(\mathcal{H}) < \gamma_{\max}(\overline{\mathcal{H}})] &= \Pr[\gamma_{\min}(\mathcal{H}) < \gamma_{\max}(\overline{\mathcal{H}})] \Pr[\mathcal{A}(k) = \mathcal{H} \mid \gamma_{\min}(\mathcal{H}) < \gamma_{\max}(\overline{\mathcal{H}})] \\ &= 0 \end{aligned} \quad (16)$$

According to our voting rule, the accepted paper set $\mathcal{A}(k)$ equals to \mathcal{H} is 1 conditioned on $\gamma_{\min}(\mathcal{H}) > \gamma_{\max}(\overline{\mathcal{H}})$. Thus,

$$\begin{aligned} \Pr[\mathcal{A}(k) = \mathcal{H}, \gamma_{\min}(\mathcal{H}) > \gamma_{\max}(\overline{\mathcal{H}})] &= \Pr[\gamma_{\min}(\mathcal{H}) > \gamma_{\max}(\overline{\mathcal{H}})] \Pr[\mathcal{A}(k) = \mathcal{H} \mid \gamma_{\min}(\mathcal{H}) > \gamma_{\max}(\overline{\mathcal{H}})] \\ &= \Pr[\gamma_{\min}(\mathcal{H}) > \gamma_{\max}(\overline{\mathcal{H}})]. \end{aligned} \quad (17)$$

Note that $\gamma_i = \sum_{j=1}^n S_j^i / n, \forall i$. Since the scores $S_j^i, \forall i, j$, are independent random variables, thus the average scores $\gamma_1, \dots, \gamma_N$ are also independent random variables. Based on this fact, we derive the analytical expression of $\Pr[\gamma_{\min}(\mathcal{H}) > \gamma_{\max}(\overline{\mathcal{H}})]$ as

$$\begin{aligned} \Pr[\gamma_{\min}(\mathcal{H}) > \gamma_{\max}(\overline{\mathcal{H}})] &= \sum_{\ell=n}^{nm} \Pr[\gamma_{\min}(\mathcal{H}) = \ell/n] \Pr[\gamma_{\max}(\overline{\mathcal{H}}) < \ell/n] \\ &= \sum_{\ell=n}^{nm} \left(\prod_{i \in \mathcal{I}} \Pr[\gamma_i \leq \ell/n] - \prod_{i \in \mathcal{I}} \Pr[\gamma_i \leq (\ell-1)/n] \right) \prod_{j \in \overline{\mathcal{I}}} \Pr[\gamma_j \leq (\ell-1)/n]. \end{aligned} \quad (18)$$

The remaining task is to derive the last term of Eq. (15). When $\gamma_{\min}(\mathcal{H}) = \gamma_{\max}(\overline{\mathcal{H}})$ occurs, tie breaking will be performed on the set of papers with the average score equal to $\gamma_{\min}(\mathcal{H})$. Let us provide some notations first. Let $\mathcal{F} = \{i \mid \gamma_i = \gamma_{\min}(\mathcal{H}), i \in \mathcal{I}\}$ be the index set of the papers that belong to set \mathcal{H} and with average scores equal to minimum average score of \mathcal{H} . Let $\mathcal{G} = \{i \mid \gamma_i = \gamma_{\min}(\mathcal{H}), i \in \overline{\mathcal{I}}\}$ be the index set of the papers that belong to set $\overline{\mathcal{H}}$ and with average scores equal to the minimum average score of \mathcal{H} . Thus, tie breaking will perform on the papers with index set $\mathcal{F} \cup \mathcal{G}$, from which only $|\mathcal{F}|$ papers will be selected for acceptance. By enumerating all possible tie breaking paper sets, we can divide the the last term of Eq. (15) into the following form:

$$\Pr[\mathcal{A}(k) = \mathcal{H}, \gamma_{\min}(\mathcal{H}) = \gamma_{\max}(\overline{\mathcal{H}})] = \sum_{\mathcal{F} \subseteq \mathcal{I}, \mathcal{G} \subseteq \overline{\mathcal{I}}, \mathcal{F}, \mathcal{G} \neq \emptyset} \Pr[\mathcal{F} \cup \mathcal{G}] \Pr[\mathcal{F} \mid \mathcal{F} \cup \mathcal{G}], \quad (19)$$

where $\Pr[\mathcal{F} \cup \mathcal{G}]$ is the probability that tie breaking is performed on papers with index set $\mathcal{F} \cup \mathcal{G}$, and $\Pr[\mathcal{F} \mid \mathcal{F} \cup \mathcal{G}]$ is the conditional probability that papers with index set \mathcal{F} is selected for acceptance under the condition that tie breaking is performed on the papers with index set $\mathcal{F} \cup \mathcal{G}$. Since the tie-breaking rule is the random rule, under which we just randomly pick $|\mathcal{F}|$ papers, thus

$$\Pr[\mathcal{F} \mid \mathcal{F} \cup \mathcal{G}] = \binom{|\mathcal{F} \cup \mathcal{G}|}{|\mathcal{F}|}^{-1}. \quad (20)$$

Because the average scores $\gamma_1, \dots, \gamma_N$ are independent random variables, we can derive $\Pr[\mathcal{F} \cup \mathcal{G}]$ as:

$$\begin{aligned} \Pr[\mathcal{F} \cup \mathcal{G}] &= \Pr[\gamma_i > \gamma_{\min}(\mathcal{H}) \quad \text{for all } i \in \mathcal{I} \setminus \mathcal{F}] \times \\ &\quad \Pr[\gamma_i = \gamma_{\min}(\mathcal{H}) \quad \text{for all } i \in \mathcal{F} \cup \mathcal{G}] \times \\ &\quad \Pr[\gamma_i < \gamma_{\min}(\mathcal{H}) \quad \text{for all } i \in \overline{\mathcal{I}} \setminus \mathcal{G}] \\ &= \sum_{\ell=n}^{nm} \prod_{i \in \mathcal{I} \setminus \mathcal{F}} (1 - \Pr[\gamma_i \leq \ell/n]) \prod_{j \in \mathcal{F} \cup \mathcal{G}} \Pr[\gamma_j = \ell/n] \prod_{\kappa \in \overline{\mathcal{I}} \setminus \mathcal{G}} \Pr[\gamma_\kappa \leq (\ell-1)/n]. \end{aligned} \quad (21)$$

Combining Eq. (19) – (21) we obtain

$$\Pr[\mathcal{A}(k) = \mathcal{H}, \gamma_{\min}(\mathcal{H}) = \gamma_{\max}(\overline{\mathcal{H}})] = \sum_{\mathcal{F} \subseteq \mathcal{I}, \mathcal{G} \subseteq \overline{\mathcal{I}}, \mathcal{F}, \mathcal{G} \neq \emptyset} \sum_{\ell=n}^{nm} \binom{|\mathcal{F} \cup \mathcal{G}|}{|\mathcal{F}|}^{-1} \prod_{i \in \mathcal{I} \setminus \mathcal{F}} (1 - \Pr[\gamma_i \leq \ell/n]) \prod_{j \in \mathcal{F} \cup \mathcal{G}} \Pr[\gamma_j = \ell/n] \prod_{\kappa \in \overline{\mathcal{I}} \setminus \mathcal{G}} \Pr[\gamma_\kappa \leq (\ell - 1)/n]. \quad (22)$$

Combining Eq. (15) – (18), and (22), we obtain Eq. (14). ■

To illustrate the analytical expression of $\Pr[\mathcal{A}(k) = \{P_{i_1}, \dots, P_{i_k}\}]$, consider a simple example where $n = 1$, $m = 2$, $N = 3$ and $k = 1$. Table 1 shows the analytical expression $\Pr[\mathcal{A}(1) = \{P_{i_1}\}]$.

$\{P_{i_1}\}$	$\Pr[\mathcal{A}(1) = \{P_{i_1}\}]$
$\{P_1\}$	$\Pr[\gamma_1 = 2]\Pr[\gamma_2 = 1]\Pr[\gamma_3 = 1] + \Pr[\gamma_1 = 2]\Pr[\gamma_2 = 2]\Pr[\gamma_3 = 1]/2 + \Pr[\gamma_1 = 2]\Pr[\gamma_3 = 2]\Pr[\gamma_2 = 1]/2 + \Pr[\gamma_1 = 1]\Pr[\gamma_3 = 1]\Pr[\gamma_2 = 1]/3$ $\Pr[\gamma_1 = 2]\Pr[\gamma_3 = 2]\Pr[\gamma_2 = 2]/3$
$\{P_2\}$	$\Pr[\gamma_2 = 2]\Pr[\gamma_1 = 1]\Pr[\gamma_3 = 1] + \Pr[\gamma_2 = 2]\Pr[\gamma_1 = 2]\Pr[\gamma_3 = 1]/2 + \Pr[\gamma_2 = 2]\Pr[\gamma_3 = 2]\Pr[\gamma_1 = 1]/2 + \Pr[\gamma_2 = 1]\Pr[\gamma_3 = 1]\Pr[\gamma_1 = 1]/3$ $\Pr[\gamma_2 = 2]\Pr[\gamma_3 = 2]\Pr[\gamma_1 = 2]/3$
$\{P_3\}$	$\Pr[\gamma_3 = 2]\Pr[\gamma_2 = 1]\Pr[\gamma_1 = 1] + \Pr[\gamma_3 = 2]\Pr[\gamma_2 = 2]\Pr[\gamma_1 = 1]/2 + \Pr[\gamma_3 = 2]\Pr[\gamma_1 = 2]\Pr[\gamma_2 = 1]/2 + \Pr[\gamma_3 = 1]\Pr[\gamma_1 = 1]\Pr[\gamma_2 = 1]/3$ $\Pr[\gamma_3 = 2]\Pr[\gamma_1 = 2]\Pr[\gamma_2 = 2]/3$

Table 1: Examples of the analytical expression of $\Pr[\mathcal{A}(1) = \{P_{i_1}\}]$, where $i_1 = 1, \dots, 3$.

Up to now, we have derived the probability of a specific set of accepted papers. Let us derive the pmf of a general set of k papers:

$$\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = i], \quad \text{for } i = 0, 1, \dots, k,$$

which is shown in the following theorem.

Theorem 2 *The pmf of $|\mathcal{A}^I(k) \cap \mathcal{A}(k)|$ is:*

$$\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = i] = \sum_{\substack{\mathcal{F} \subseteq \mathcal{A}^I(k) \\ |\mathcal{F}|=i}} \sum_{\substack{\mathcal{G} \subseteq \overline{\mathcal{A}^I(k)} \\ |\mathcal{G}|=k-i}} \Pr[\mathcal{A}(k) = \mathcal{F} \cup \mathcal{G}], \quad \text{for all } i = 0, 1, \dots, k, \quad (23)$$

where $\overline{\mathcal{A}^I(k)} = \{P_1, \dots, P_N\} / \mathcal{A}^I(k)$ is the complement of $\mathcal{A}^I(k)$ and $\Pr[\mathcal{A}(k) = \mathcal{F} \cup \mathcal{G}]$ is given in Eq. (14).

Proof: The paper set $\mathcal{A}(k)$ can be divided into two disjoint subsets of which one is $\mathcal{A}(k) \cap \mathcal{A}^I(k)$ and the other one is $\mathcal{A}(k) \cap \overline{\mathcal{A}^I(k)}$. Note that we have derived the pmf that a specific set of accepted papers in Eq. (14). Then by enumerating subsets of $\mathcal{A}(k)$ with cardinality i and all the subsets of $\overline{\mathcal{A}^I(k)}$ with cardinality $k - i$ we can obtain probability $\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = i]$, or

$$\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = i] = \sum_{\substack{\mathcal{F} \subseteq \mathcal{A}^I(k) \\ |\mathcal{F}|=i}} \sum_{\substack{\mathcal{G} \subseteq \overline{\mathcal{A}^I(k)} \\ |\mathcal{G}|=k-i}} \Pr[\mathcal{A}(k) = \mathcal{F} \cup \mathcal{G}], \quad \text{for all } i = 0, 1, \dots, k,$$

where $\Pr[\mathcal{A}(k) = \mathcal{F} \cup \mathcal{G}]$ is given in Eq. (14). ■

i	$\Pr[\mathcal{A}^I(1) \cap \mathcal{A}(1) = i]$
1	$\Pr[\gamma_1 = 2]\Pr[\gamma_2 = 1]\Pr[\gamma_3 = 1] + \Pr[\gamma_1 = 2]\Pr[\gamma_2 = 2]\Pr[\gamma_3 = 1]/2 +$ $\Pr[\gamma_1 = 2]\Pr[\gamma_3 = 2]\Pr[\gamma_2 = 1]/2 + \Pr[\gamma_1 = 1]\Pr[\gamma_3 = 1]\Pr[\gamma_2 = 1]/3$ $\Pr[\gamma_1 = 2]\Pr[\gamma_3 = 2]\Pr[\gamma_2 = 2]/3$
0	$\Pr[\gamma_2 = 2]\Pr[\gamma_1 = 1]\Pr[\gamma_3 = 1] + \Pr[\gamma_3 = 2]\Pr[\gamma_2 = 1]\Pr[\gamma_1 = 1] +$ $\Pr[\gamma_2 = 2]\Pr[\gamma_1 = 2]\Pr[\gamma_3 = 1]/2 + \Pr[\gamma_3 = 2]\Pr[\gamma_1 = 2]\Pr[\gamma_2 = 1]/2$ $\Pr[\gamma_2 = 2]\Pr[\gamma_3 = 2]\Pr[\gamma_1 = 1] + \Pr[\gamma_2 = 1]\Pr[\gamma_3 = 1]\Pr[\gamma_1 = 1] \times 2/3$ $\Pr[\gamma_2 = 2]\Pr[\gamma_3 = 2]\Pr[\gamma_1 = 2] \times 2/3$

Table 2: Examples of the analytical expression of $\Pr[\mathcal{A}^I(1) \cap \mathcal{A}(1) = i]$, where $i = 0, 1$.

To illustrate the analytical expression of $\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = i]$, let us consider the same example with Table 1. Table 2 shows the analytical expression $\Pr[\mathcal{A}^I(1) \cap \mathcal{A}(1) = i]$, where $i = 0, 1$.

Now we have derived the analytical expression of the pmf of $|\mathcal{A}^I(k) \cap \mathcal{A}(k)|$, then it is easy to obtain the analytical expression of $E[|\mathcal{A}^I(k) \cap \mathcal{A}(k)|]$ and $\text{Var}[|\mathcal{A}^I(k) \cap \mathcal{A}(k)|]$ according to the definition of expectation and variance.

Examining Eq. (23), we can see that $\Pr[\mathcal{A}(k) = \mathcal{F} \cup \mathcal{G}]$ is an essential part of the analytical expression of the pmf. The analytical expression of $\Pr[\mathcal{A}(k) = \mathcal{F} \cup \mathcal{G}]$ is given in Eq. (14), which is quite complicated, and these analytical expressions cannot be reduced to a simple form. Thus, it is not easy to use these analytical results to gain some insight of a conference recommendation system. An alternative way is to compute the numerical results of the pmf of $|\mathcal{A}^I(k) \cap \mathcal{A}(k)|$ based on the analytical expressions in Eq. (23). After we obtain the numerical results of the pmf of $|\mathcal{A}^I(k) \cap \mathcal{A}(k)|$, we can then compute its expectation and variance. Unfortunately, computing the numerical results of Eq. (23) is computationally expensive, which is shown in the following theorem.

Theorem 3 *The computational complexity in calculating the numerical results of the pmf of $|\mathcal{A}^I(k) \cap \mathcal{A}(k)|$ based on Eq. (23) is exponential, or $\Theta(2^N)$.*

Proof: Examining Eq. (23), we can see that the calculation of $\Pr[\mathcal{A}(k) = \mathcal{F} \cup \mathcal{G}]$ is the core part on the calculation of the pmf of $|\mathcal{A}^I(k) \cap \mathcal{A}(k)|$. Assume the running time of calculating $\Pr[\mathcal{A}(k) = \mathcal{F} \cup \mathcal{G}]$ is t , then by Equation (23), the running time of calculating the pmf is

$$\Theta\left(\sum_{i=0}^k \binom{k}{i} \binom{N-k}{k-i} t\right) = \Theta\left(\binom{N}{k} t\right). \quad (24)$$

In the following we analyze the running time of calculating the numerical result of $\Pr[\mathcal{A}(k) = \mathcal{F} \cup \mathcal{G}]$ based on its analytical expression derived by Eq. (14). Examining Eq. (14), we can see that there are two basic computations of Eq. (14), of which the first one is

$$\left(\prod_{i \in \mathcal{I}} \Pr\left[\gamma_i \leq \frac{\ell}{n}\right] - \prod_{i \in \mathcal{I}} \Pr\left[\gamma_i \leq \frac{\ell-1}{n}\right]\right) \prod_{j \in \overline{\mathcal{I}}} \Pr\left[\gamma_j \leq \frac{\ell-1}{n}\right],$$

let us assume the running time of calculating this basic part is t_1 . The second basic computation is

$$\prod_{i \in \mathcal{I} \setminus \mathcal{F}} \left(1 - \Pr\left[\gamma_i \leq \frac{\ell}{n}\right]\right) \prod_{j \in \mathcal{F} \cup \mathcal{G}} \Pr\left[\gamma_j \leq \frac{\ell}{n}\right] \prod_{\kappa \in \overline{\mathcal{I}} \setminus \mathcal{G}} \Pr\left[\gamma_\kappa \leq \frac{\ell-1}{n}\right],$$

let us assume the running time of calculating this basic part is t_2 . The running time of computing $\Pr[\mathcal{A}(k) = \mathcal{F} \cup \mathcal{G}]$ by Eq. (14) is

$$\begin{aligned}
t &= \Theta(npt_1 + (2^k - 1)(2^{N-k} - 1)npt_2) \\
&= \Theta(npt_1 + 2^{k-1}2^{N-k-1}npt_2) \\
&= \Theta(npt_1 + 2^{N-2}npt_2) = \Theta(2^N npt_2/4).
\end{aligned}$$

By letting $t = \Theta(2^N n p t_2 / 4)$ in Eq. (24), we can obtain the result stated in this theorem. ■

To illustrate the complexity, consider the number of mathematical operations (i.e., addition, subtraction, multiplication and division) that we need in the computation of the numerical results of the pmf of $|\mathcal{A}^I(k) \cap \mathcal{A}(k)|$. Table 3 illustrates the computational complexity of three conferences: Recsys'11, Sigcomm'11, WWW'11 and IEEE Infocom'11. Since the expectation and variance are based on these operations, so their computational complexity are also exponential.

Conference	N	k	Acceptance %	Complexity
ACM Recsys'11	110	22	20%	$\sim 2^{161}$
ACM Sigcomm'11	223	32	13%	$\sim 2^{312}$
WWW'11	658	81	12%	$\sim 2^{902}$
IEEE Infocom'11	1823	291	16%	$\sim 2^{2593}$

Table 3: Examples of computational complexity

In summary, we have the following conclusions in analyzing the above conference recommendation system:

- We can analytically derive the pmf of:

$$\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = i], \quad \text{for } i = 0, 1, \dots, k,$$

- The analytical expression is complex and it is not easy to obtain insights of the underlying recommendation system.
- Computing the numerical results based on these analytical results is computational expensive.

3.2 Derivation for the General Case

For the general case, we can derive the analytical expressions of the pmf, expectation and variance of $|\mathcal{A}^I(k) \cap \mathcal{A}(k)|$ with similar methods used in the special case. The analytical expressions for the general case will have a similar form compared with the analytical expressions derived in the special case. Furthermore, it is reasonable to expect that for the general case, there can be different types of paper and that reviewers are not homogeneous (e.g, they may have different topics preference). Also, tie breaking rules will be more complicated than the random rule. Hence we expect the analytical expressions for the general case will be more complicated. Thus, one may not easily obtain insight by examining the analytical expression of the general case. Instead, let us focus on finding a practical approach to solve the general case, and that it should be *computational inexpensive* to obtain numerical results of:

$$\Pr[|\mathcal{A}^I(k) \cap \mathcal{A}(k)| = i], \quad \text{for } i = 0, 1, \dots, k.$$

In the following section, we present this practical approach, and to show that not only we can have a computational efficient approach to compute all probability measures, but more importantly, provides performance guarantees on our operation.

4 Randomized Algorithm

In this section, we present a randomized algorithm to evaluate the pmf, expectation and variance of $|\mathcal{A}^I(k) \cap \mathcal{A}(k)|$. Our randomized algorithm is computationally efficient with performance guarantee. We will use the following notations to describe our algorithm. We define $I(k) = |\mathcal{A}^I(k) \cap \mathcal{A}(k)|$. Let $\widehat{\Pr}[I(k) = i]$, $\widehat{E}[I(k)]$ and $\widehat{\text{Var}}[I(k)]$ denote the approximate value of $\Pr[I(k) = i]$, $E[I(k)]$ and $\text{Var}[I(k)]$ respectively. We first present the algorithm, then show its performance guarantee.

4.1 Randomized Algorithm

Our algorithm is stated in Algorithm 1. The main idea of this randomized algorithm is that we first approximate the pmf $\Pr[I(k) = i]$, then we use the approximate value $\hat{\Pr}[I(k) = i]$ to compute $\hat{E}[I(k)]$ and $\hat{\text{Var}}[I(k)]$. We can

Algorithm 1 : Randomized Algorithm

```

1: for all  $i = 0, \dots, k, \ell_i \leftarrow 0$ 
2: for  $j = 1$  to  $K$  do
3:   for all  $i = 1, \dots, N$ , generate the intrinsic quality for paper  $P_i$  by simulating the paper submission process according to a self-selectivity type.
4:   produce the intrinsic top- $k$  paper set  $\mathcal{A}^I(k)$ .
5:   assign papers to reviewers, generate the degree of criticality based on the paper-reviewer matching degree.
6:   for  $i = 1, \dots, N$ , generate score set  $\mathcal{S}^i$  for paper  $P_i$  by simulating the scoring process.
7:   simulate the decision making process, i.e., applying the voting rule  $\mathcal{V}$  and the tie breaking rule  $\mathcal{T}$  to produce the set  $\mathcal{A}(k)$  based on the score sets  $\{\mathcal{S}^1, \dots, \mathcal{S}^N\}$ 
8:   if the cardinality of the intersection of  $\mathcal{A}^I(k)$  and  $\mathcal{A}(k)$  is equal to  $i$ , then  $\ell_i \leftarrow \ell_i + 1$ .
9: end for
10: for all  $i = 0, \dots, k, \hat{\Pr}[I(k) = i] \leftarrow \ell_i / K$ 
11:  $\hat{E}[I(k)] \leftarrow \sum_{i=0}^k i \hat{\Pr}[I(k) = i]$ 
12:  $\hat{\text{Var}}[I(k)] \leftarrow \sum_{i=0}^k (i - \hat{E}[I(k)])^2 \hat{\Pr}[I(k) = i]$ 

```

state two properties of this algorithm. The first one is its running time complexity and the other one is its theoretical performance guarantee. The following theorem states its running time complexity.

Theorem 4 *The computational complexity of our randomized algorithm is $\Theta(KN \log N)$, where K is the number of simulation round and N is the number of submitted papers.*

Proof: We prove this theorem by examining the complexity of each step of our Algorithm 1. The complexity of step 10 – 12 and 1 are the same, say $\Theta(k)$. The complexity of step 3 – 8 are $\Theta(N)$, $\Theta(N \log N)$, $\Theta(\sum_{i=1}^N n_i)$, $\Theta(\sum_{i=1}^N n_i)$, $\Theta(N \log N)$, and $\Theta(1)$ respectively. Since each reviewer only review a small subset of the submitted papers, namely $n_i \ll N$, thus we have $\Theta(\sum_{i=1}^N n_i) = \Theta(N)$. By adding the complexity of step 1 – 12 up we could obtain the theorem. ■

The remaining technical issue is how to set the parameter K . Specifically, how many simulation rounds K can produce a good approximation of the pmf, expectation and variance? Let us proceed to answer this question by deriving the theoretical performance guarantee for Algorithm 1.

4.2 Theoretical Performance Guarantee

First, we derive a loose bound on the number of simulation rounds K needed but have good performance guarantee. Then we show how one can have a tight bound on K , and tradeoff between K and its performance guarantee. The following theorem states the loose bound on K .

Theorem 5 *When the following condition holds:*

$$K \geq \max_{\substack{i=0,1,\dots,k, \\ \Pr[I(k)=i] \neq 0}} \frac{3 \ln(2(k+1)/\delta)}{\Pr[I(k) = i] \epsilon^2}, \quad (25)$$

then Algorithm 1 guarantees:

$$\begin{aligned}
\left| \hat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| &\leq \epsilon \Pr[I(k) = i], \quad \forall i = 0, \dots, k, \\
\left| \hat{E}[I(k)] - E[I(k)] \right| &\leq \epsilon E[I(k)], \\
\left| \hat{\text{Var}}[I(k)] - \text{Var}[I(k)] \right| &\leq \epsilon(1 + \epsilon) \text{Var}[I(k)],
\end{aligned}$$

with probability at least $1 - \delta$.

Proof: By applying Lemma 4, we obtain that

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \Pr[I(k) = i],$$

holds for all $x = 0, 1, \dots, k$, with probability at least $1 - \delta$. By applying Lemma 5 and 4 we have

$$\left| \widehat{E}[I(k)] - E[I(k)] \right| \leq \epsilon E[I(k)]$$

holds with probability at least $1 - \delta$. By applying Lemma 7 and 4 we have

$$\left| \widehat{\text{Var}}[I(k)] - \text{Var}[I(k)] \right| \leq \epsilon(1 + \epsilon)\text{Var}[I(k)]$$

holds with probability at least $1 - \delta$. Please refer appendix for the proofs of these lemmas. ■

When one examines the Inequality (25), one can see that the bound of K is useful when $\Pr[I(k) = i]$ is not small for $i = 0, 1, \dots, k$. Consider the case when $\Pr[I(k) = i] \leq 2^{-N}$ for some $i = 0, 1, \dots, k$, then we have $K \geq 2^N$. For such cases, K is too large. In the following theorem, we show a tight bound on K .

Theorem 6 When $K \geq 3 \ln(2(k+1)/\delta)/\epsilon^2$, Algorithm 1 guarantees:

$$\begin{aligned} \left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| &\leq \epsilon \sqrt{\Pr[I(k) = i]}, \quad \forall i = 0, \dots, k, \\ \left| \widehat{E}[I(k)] - E[I(k)] \right| &\leq \epsilon \sqrt{k(k+1)E[I(k)]/2} \\ \left| \widehat{\text{Var}}[I(k)] - \text{Var}[I(k)] \right| &\leq \epsilon(k+1) \left(\epsilon \text{Var}[I(k)] + \sqrt{(2k+1)\text{Var}[I(k)]/6} \right) \end{aligned}$$

with probability at least $1 - \delta$.

Proof: By applying Lemma 9, we obtain that

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \sqrt{\Pr[I(k) = i]}$$

holds for all $i = 0, 1, \dots, k$, with probability at least $1 - \delta$. By applying Lemma 10, and 9 we have

$$\left| \widehat{E}[I(k)] - E[I(k)] \right| \leq \epsilon \sqrt{k(k+1)E[I(k)]/2}$$

holds with probability $1 - \delta$. By applying Lemma 12 and 9 we have

$$\left| \widehat{\text{Var}}[I(k)] - \text{Var}[I(k)] \right| \leq \epsilon(k+1) \left(\epsilon \text{Var}[I(k)] + \sqrt{(2k+1)\text{Var}[I(k)]/6} \right)$$

hold with probability at least $1 - \delta$. Detailed proof of these lemmas are in the appendix. ■

5 Evaluation of Peer Review System

In this section we evaluate the accuracy of conference recommendation systems. We consider a conference with two hundred submissions, or $N = 200$, and only $k = 30$ submissions will be accepted. In other words, a 15% acceptance rate. In consistent with realistic conference review systems, we set $m = 5$, or the rating set is $\{1, \dots, 5\}$. The simulation rounds K in Algorithm 1 is set to 10^9 . In the following, we start our evaluation from a simple case, then we extend it step by step and evaluate the impact of various factors on the overall accuracy in the final recommendation.

5.1 Probability distribution, expectation and variance of $|\mathcal{A}^I(k) \cap \mathcal{A}(k)|$

Here we consider a homogenous conference recommendation system, in which papers and reviewers are homogeneous and each paper is reviewed by the *same number* of reviewers, or $n_i = n, \forall i$. Specifically, each reviewer is non-biased with the same critical degree, or $c_j^i = c, \forall i, j$. Hence the function $\sigma(c_j^i), \forall i, j$, within Equation (6) has the same value $\sigma(c)$ and further we set $\sigma(c) = 1$. We select one type of *self-selectivity*, say medium *self-selectivity* specified by Equation (2), to study here. The voting rule \mathcal{V} is the *average score rule* and the tie breaking rule is the *least variance rule* that selects the paper with the least variance. If there is still tie, paper is randomly selected.

Definition 1 The reviewing workload of a conference recommendation system is the sum of all reviews, or

$$W = \sum_{i=1}^N n_i.$$

Definition 2 Let I_i be the random variable indicating

$$I_i = |\mathcal{A}^I(i) \cap \mathcal{A}(k)|. \quad (26)$$

In other words, if $i = 1$, I_1 reflects the event that the best paper is accepted by this conference recommendation system. When $i = 5$, I_5 reflects the event that the top five submitted papers are finally accepted.

The numerical results of the pmf of $|\mathcal{A}^I(30) \cap \mathcal{A}(30)|$ are shown in Fig. 1. The numerical results of the expectation and variance of I_1, I_5, I_{10} and I_{30} are shown in Table 4.

In Fig. 1, the horizontal axis represents the number of top 30 papers that got finally accepted, or $|\mathcal{A}^I(30) \cap \mathcal{A}(30)|$. The vertical axis shows the corresponding probability. From Fig. 1 we could see that when we increase the number of reviews per paper, or n , the probability mass function shifts toward the right. In other words, the more reviews each paper received, the higher the accuracy of the conference recommendation system. From Table 4, we have the following observations. When each paper is reviewed by three reviewers, approximately 19.8 papers from the top 30 papers will get accepted. It is interesting to note that the chance of accepting the best paper is invariant of the reviewing workload, since $E[I_1] = 0.98$ when $n = 3$ and it improves to 0.9999 when $n = 10$. This statement also holds for the top ten papers. From Table 4, we can see that as we increase n , we decrease the variance, which reflects that the conference recommendation system is more accurate.

Lessons learned: If reviewers are non-biased, fair and critical, and the quality of submitted papers is of medium *self-selectivity* as specified by Eq. (2), we have a pretty accurate conference recommendation system. There are number of interesting questions to explore further, i.e., are these results dependent on the distribution of intrinsic quality, or the *self-selectivity* type of papers? Are these results sensitive to any voting rule? Let us continue to explore.

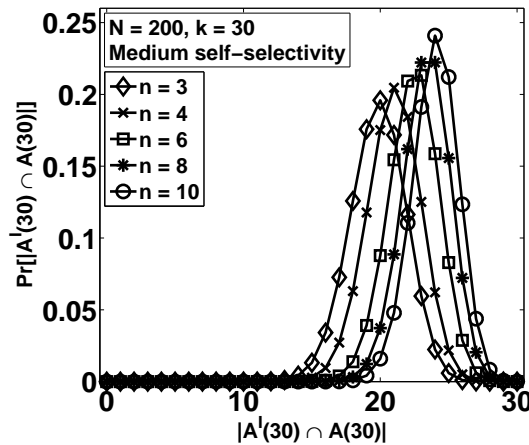


Figure 1: pmf of $|\mathcal{A}^I(30) \cap \mathcal{A}(30)|$ when $n = 3, 4, 6, 8, 10$ and papers are submitted with medium *self-selectivity*.

	$n = 3$	$n = 4$	$n = 6$	$n = 8$	$n = 10$
$E[I_1]$	0.9832	0.9931	0.9986	0.9996	0.9999
$E[I_5]$	4.6854	4.8284	4.9352	4.9723	4.9869
$E[I_{10}]$	8.7763	9.1643	9.5576	9.7446	9.8426
$E[I_{30}]$	19.8270	20.960	22.400	23.328	23.980
$\text{Var}[I_1]$	0.0165	0.0067	0.0014	0.0004	0.0001
$\text{Var}[I_5]$	0.2942	0.1701	0.0654	0.0280	0.0132
$\text{Var}[I_{10}]$	1.0793	0.7793	0.4370	0.2586	0.1609
$\text{Var}[I_{30}]$	4.1210	3.7710	3.2934	2.9580	2.7121

Table 4: Expectation and variance of I_1 , I_5 , I_{10} , and I_{30} when we increase number of reviews per paper, $n = 3, 4, 6, 8, 10$, and papers are submitted with medium *self-selectivity*.

5.2 Effect of Intrinsic Quality

Here, we explore the effect of *intrinsic quality* (or *self-selectivity*) of papers on the conference recommendation system. Specifically, we consider four representative types of *self-selectivity* of papers: high, medium, low and random *self-selectivity* specified by Eq. (1) – (4) respectively. We explore the effect of *self-selectivity* of papers on the homogeneous conference recommendation system specified in Section 5.1 with each paper receiving three reviews, or $n = 3$. We use the following notations to present our results:

H-S-S: high *self-selectivity* specified in Eq. (1).

M-S-S: medium *self-selectivity* specified in Eq. (2).

L-S-S: low *self-selectivity* specified in Eq. (3).

R-S-S: random *self-selectivity* specified in Eq. (4).

The numerical results of the pmf of $|\mathcal{A}^I(30) \cap \mathcal{A}(30)|$ is shown in Fig. 2. The numerical results of the expectation and variance of I_1 , I_5 , I_{10} and I_{30} are shown in Table 5.

In Fig. 2, the horizontal axis represents the number of top 30 papers that got finally accepted, or $|\mathcal{A}^I(30) \cap \mathcal{A}(30)|$. The vertical axis shows the corresponding probability. From Fig. 2 we could see that as the *self-selectivity* type varies in the order of high, low, medium, random *self-selectivity*, the corresponding mass probability distribution curve moves towards right. In other words, the accuracy of the conference recommendation system corresponding to the random *self-selectivity* is the highest followed by medium, low and high *self-selectivity*. From Table 5 we have the following observations. When papers are submitted with medium, low or random *self-selectivity*, around 20 papers from the top 30 papers will be accepted. It is interesting to note that the chance of accepting the best paper is invariant to these three *self-selectivity* types, since the corresponding three values of $E[I_1]$ are all around 0.98. This statement also holds for top ten papers. But when papers are submitted with high *self-selectivity*, the accuracy of the conference recommendation system is *remarkably lower* than that the other three *self-selectivity* types. Specifically, only a small number, around 13.2, of papers from the top 30 papers will be accepted. Even the best paper will get rejected with high probability, around 0.46. This statement also holds for top ten or top five papers. The variance corresponding to the high *self-selectivity* is the highest among those four, which reflects that the results of the conference recommendation system is the least accurate and more likely to depart from the expectation. Hence, for a prestigious conference, i.e., SIGCOMM, assume papers are submitted with high *self-selectivity*, and if the conference insists to have a small technical program committee with reviewers having moderate reviewing workload, say $n = 3$, the final recommendation may not be accurate.

Lessons learned: When reviewers are non-biased fair and critical, the above simple conference review system is quite accurate except when papers are of high *self-selectivity*. In that case, one may explore other means to improve the accuracy. Again, there are number of interesting questions to explore, i.e., how large the reviewing workload do we need to have an accurate recommendation?

5.3 Effect of Number of Reviews per Paper

Here we consider a homogeneous conference recommendation system specified in Section 5.1. We select one probability measure, expectation of acceptance, to study. The numerical results of $E[I_1]$, $E[I_5]$, $E[I_{10}]$ and $E[I_{30}]$ are

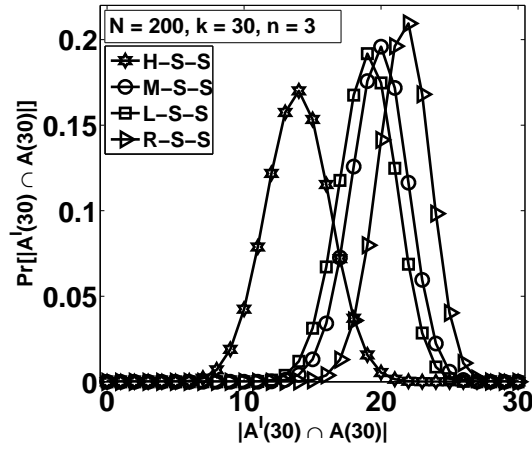


Figure 2: pmf of $|A^I(30) \cap A(30)|$ when $n = 3$ and papers are submitted with high, medium, low or random *self-selectivity*.

	<i>H-S-S</i>	<i>M-S-S</i>	<i>L-S-S</i>	<i>R-S-S</i>
$E[I_1]$	0.5401	0.9832	0.9788	0.9846
$E[I_5]$	2.6250	4.6950	4.6082	4.7599
$E[I_{10}]$	5.0687	8.7736	8.5162	9.0810
$E[I_{30}]$	13.2258	19.8270	18.9798	21.5821
$\text{Var}[I_1]$	0.2437	0.0165	0.0208	0.0151
$\text{Var}[I_5]$	1.2236	0.2942	0.3698	0.2314
$\text{Var}[I_{10}]$	2.3933	1.0793	1.2491	0.8275
$\text{Var}[I_{30}]$	5.4420	4.1210	4.2888	3.5473

Table 5: Expectation and Variance of I_1 , I_5 , I_{10} , and I_{30} when $n = 3$ and papers are submitted with high, medium, low or random *self-selectivity*.

shown in Fig. 3.

In Fig. 3, the horizontal axis represents the number of reviews per paper, or n . The vertical axis shows the corresponding expectation. From Fig. 3, we have the following observations. When we increase the number of reviews per paper, or n , the expectation increased, which reflects the improvement in accuracy of the conference recommendation system. As the *self-selectivity* type varies in the order of high, low, medium, random *self-selectivity*, the expectation curve shifts toward up. In other words, the accuracy corresponding to the random *self-selectivity* is the highest followed by medium, low, high *self-selectivity*. It is interesting to note that the best paper is invariant of the workload, except for papers with high *self-selectivity*. This statement also holds for the top five or ten papers. When papers are highly *self-selective*, the accuracy of the conference recommendation system is remarkably lower than the other three *self-selectivity* types. Especially, when the reviewing workload is low, say $n = 3$, with less than 15 papers from the top 30 papers will get accepted. This holds even for the best paper, which only has a probability of less than 0.6 of being accepted when $n = 3$. Same can be said for the top five or top ten papers. In closing, for papers with high *self-selectivity*, and we may have to increase the reviewing workload to at least $n \geq 7$ such that we have a strong guarantee that the best paper will be accepted.

Lessons learned: If reviewers are non-biased and fair, using the above simple conference review system achieves relatively high accuracy except when papers are of high *self-selectivity*. In that case, we have to increase the workload to $n \geq 7$ to improve the system. Again, there are number of interesting questions to explore, i.e., is increasing the workload the only way to improve the conference recommendation system? Can we improve it by using different voting rule or tie breaking rule?

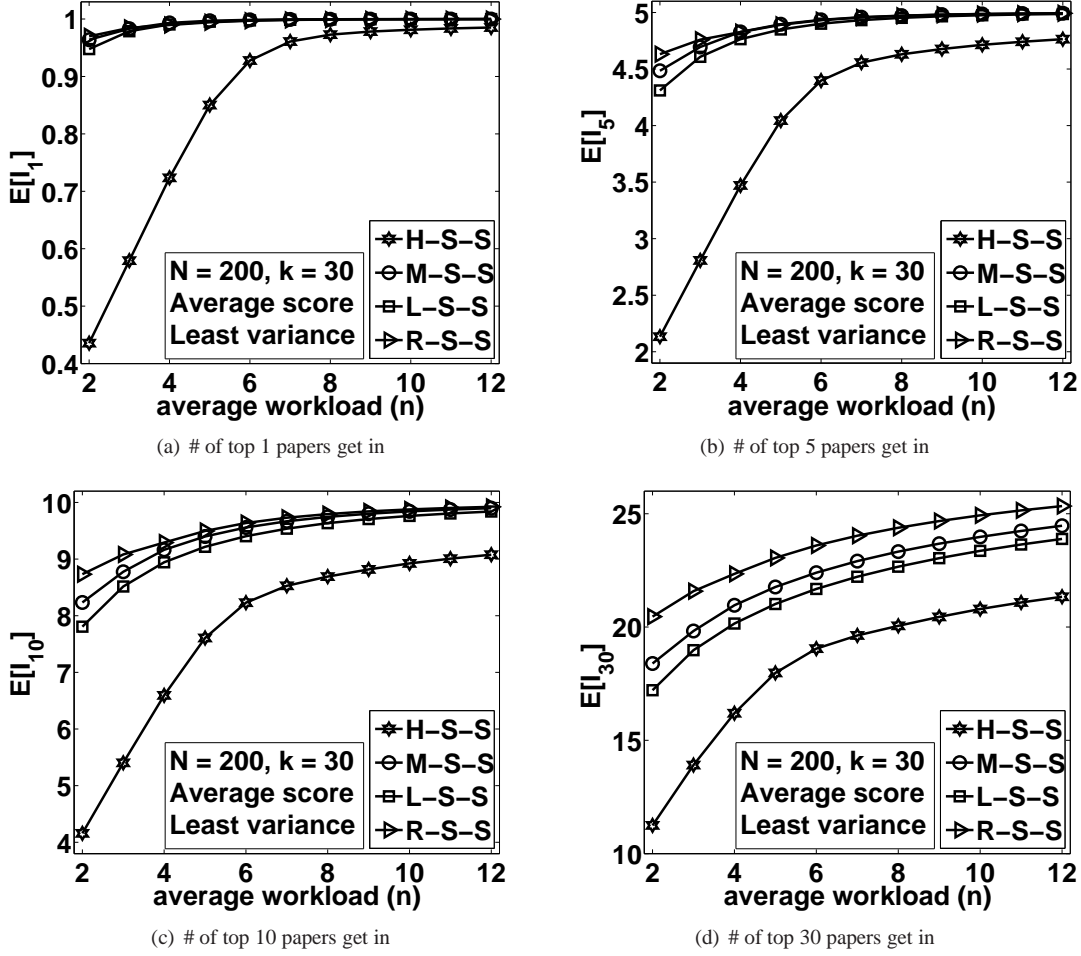


Figure 3: Expectation of I_1 , I_5 , I_{10} , and I_{30} when we vary number of reviews per paper $n = 2, \dots, 12$, with papers are from high, low, medium or random *self-selectivity*.

5.4 Effect of Voting Rules

In this section we explore the effect of voting rules on the accuracy of conference recommendation system. Specifically, we will evaluate the performance of the following three representative voting rules:

Average score rule (\mathcal{V}_{as}): specified by

$$\gamma_i = \sum_{S \in \mathcal{S}^i} S / |\mathcal{S}^i|. \quad (27)$$

Eliminate the highest & lowest score rule (\mathcal{V}_{ehl}): eliminate the highest and lowest score of each paper, and calculate the average score of each paper by the remaining scores, or

$$\gamma_i = \left(\sum_{S \in \mathcal{S}^i} S - \max \{ \mathcal{S}^i \} - \min \{ \mathcal{S}^i \} \right) / (|\mathcal{S}^i| - 2). \quad (28)$$

Punish low scores rule (\mathcal{V}_{pl}): punish the low scores. Specifically, a low score, or 1, brings the an extra punishment of decreasing its score by η , or

$$\gamma_i = \sum_{S \in \mathcal{S}^i} S / |\mathcal{S}^i| - \eta |\{S \mid S = 1, S \in \mathcal{S}^i\}|. \quad (29)$$

Let us set the punishment η to be 0.5 throughout this paper.

We use the *least variance* rule for tie breaking and we choose expectation as our performance measure. We evaluate the accuracy of these three voting rules on the conference recommendation system specified in Section 5.1. The numerical results of expectation of I_{30} are shown in Fig. 4. The numerical results of expectation of I_1 and I_5 for

	\mathcal{V}_{as}	\mathcal{V}_{ehl}	\mathcal{V}_{pl}
High self-selectivity			
$E[I_1]$	0.5793	0.5793	0.5793
$E[I_5]$	2.8078	2.8078	2.8077
$E[I_{10}]$	5.4028	5.4026	5.4026
Medium self-selectivity			
$E[I_1]$	0.9832	0.9911	0.9832
$E[I_5]$	4.6950	4.7733	4.6950
$E[I_{10}]$	8.7763	8.9641	8.7763
Low self-selectivity			
$E[I_1]$	0.9788	0.9741	0.9746
$E[I_5]$	4.6082	4.5616	4.5691
$E[I_{10}]$	8.5162	8.4032	8.4297
Random self-selectivity			
$E[I_1]$	0.9846	0.9873	0.9846
$E[I_5]$	4.7599	4.7937	4.7599
$E[I_{10}]$	9.0810	9.1791	9.0810

Table 6: Expectation of $E[I_1]$, $E[I_5]$, and $E[I_{10}]$ for three voting rules: \mathcal{V}_{as} , \mathcal{V}_{ehl} , and \mathcal{V}_{pl} when $n = 3$

high *self-selectivity* submissions are shown in Fig. 5. The numerical results of $E[I_1]$, $E[I_5]$, and $E[I_{10}]$ when $n = 3$ are shown in Table 6.

In Fig. 4 and 5 the horizontal axis represents the number of reviews per paper, or n . The vertical axis shows corresponding expectation. From Fig. 4, we have the following observations. When we increase n , the expectation $E[I_{30}]$ corresponding to each voting rule increased, which reflects that the accuracy of the conference recommendation system increased. From Fig. 4(a), we see that when submitted papers are of high *self-selectivity*, the expectation curves overlapped together. In other words, for high *self-selectivity* papers, these three voting rules are similar. From Fig. 4(b) we see that for submitted papers with medium *self-selectivity*, the *average score rule* and *punish low scores rule* have the same degree of accuracy, and the *eliminate highest and lowest score rule* has slightly higher accuracy than the other two voting rules. This statement also holds for the random *self-selectivity* submissions as shown in Fig. 4(d). From Fig. 4(c), we see that when submitted papers are of low *self-selectivity*, all three voting rules have nearly the same accuracy but the *punish low scores rule* has slightly lower accuracy. From Table 6, we observe that when $n = 3$, the chance of accepting the best paper is invariant of these voting rules, unless when submitted papers are of high *selectivity*, since $E[I_1]$ is around 0.98 for medium, low and random *self-selectivity* papers. This statement also holds for the top five and the top ten papers. When papers are submitted with high *self-selectivity*, the chance of accepting the best paper is low, in fact, it is less than 0.6. The same statements holds for the top five and top ten papers. From Fig. 5, we see that when papers submitted with high *self-selectivity*, the expectation curves overlapped together. And for each voting rule we have to increase the reviewing workload to at least seven such that we have a strong guarantee that the best paper or the top five papers will get accepted.

Lessons learned: These three voting rules have comparable accuracy, no rule can outperform others remarkably. Thus the improvement of conference recommendation system by voting rules is limited. For each voting rule, we still have to increase the average workload to at least seven to have a strong guarantee that the best paper will get in. Again, there are number of interesting questions to explore, i.e., how about the tie-breaking rules?

5.5 Effect of Tie Breaking Rules

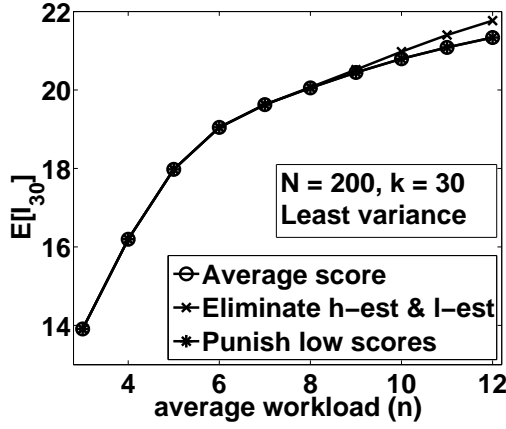
In this section we explore the effect of tie breaking rules on conference recommendation system. Specifically, we evaluate the performance of the following four representative tie breaking rules:

Least variance (\mathcal{T}_{var}): select one with least variance, if there is still tie, perform random selection.

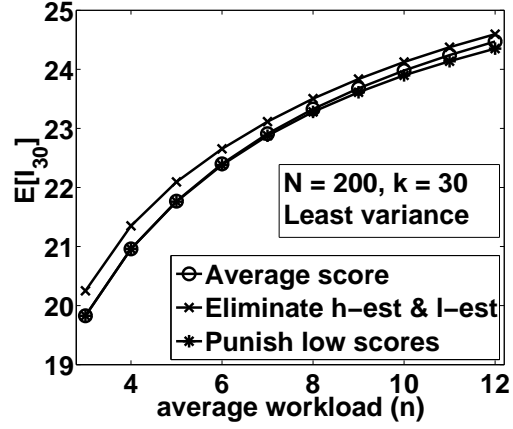
Largest max score (\mathcal{T}_{maxs}): select one with the largest max score, if there is still tie, perform random selection.

Largest min score (\mathcal{T}_{mins}): select one with the largest min score, if there is still tie, perform random selection.

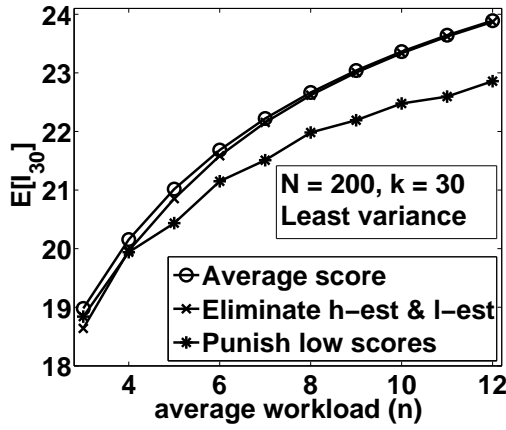
Largest medium score (\mathcal{T}_{meds}): select one with the largest medium score, if there is still tie, perform random selection.



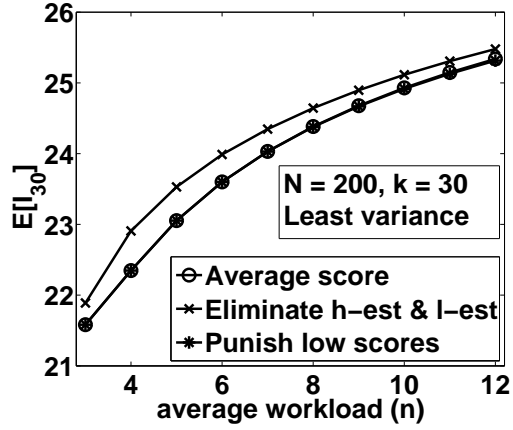
(a) High self-selectivity



(b) Medium self-selectivity

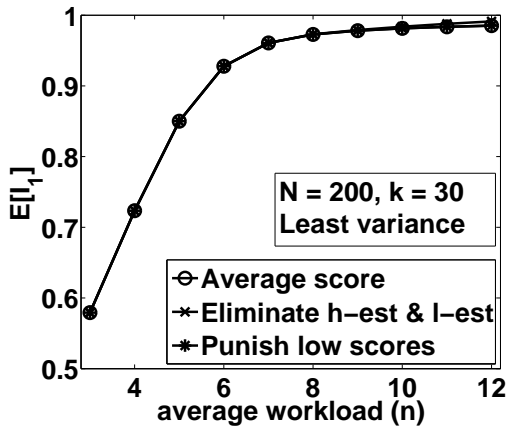


(c) Low self-selectivity

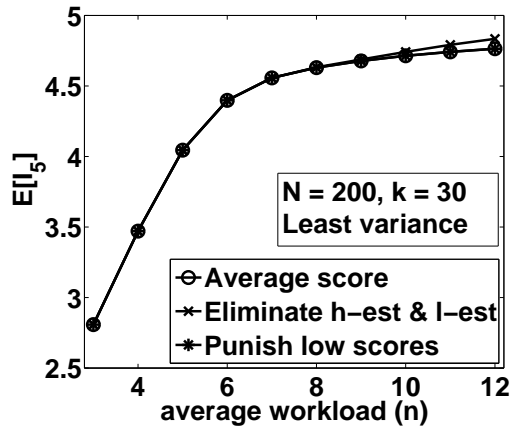


(d) Random self-selectivity

Figure 4: $E[I_{30}]$ for three voting rules: \mathcal{V}_{as} , \mathcal{V}_{ehl} , and \mathcal{V}_{pl} .



(a) # of top 1 paper get in



(b) # of top 5 papers get in

Figure 5: Expectation of $E[I_1]$, $E[I_5]$ for three voting rules: \mathcal{V}_{as} , \mathcal{V}_{ehl} , and \mathcal{V}_{pl} for high self-selectivity submission.

	\mathcal{T}_{var}	\mathcal{T}_{maxs}	\mathcal{T}_{mins}	\mathcal{T}_{meds}
High self-selectivity				
$E[I_1]$	0.5793	0.5793	0.5793	0.5793
$E[I_5]$	2.8078	2.8077	2.8077	2.8078
$E[I_{10}]$	5.4026	5.4026	5.4026	5.4026
Medium self-selectivity				
$E[I_1]$	0.9832	0.9880	0.9832	0.9907
$E[I_5]$	4.6951	4.7514	4.6951	4.7755
$E[I_{10}]$	8.7764	8.9311	8.7765	8.9831
Low self-selectivity				
$E[I_1]$	0.9788	0.9787	0.9788	0.9886
$E[I_5]$	4.6083	4.5989	4.6074	4.6040
$E[I_{10}]$	8.5164	8.4723	8.5102	8.3990
Random self-selectivity				
$E[I_1]$	0.9846	0.9858	0.9846	0.9872
$E[I_5]$	4.7599	4.7747	4.7599	4.7926
$E[I_{10}]$	9.0810	9.1239	9.0810	9.1761

Table 7: Expectation of $E[I_1]$, $E[I_5]$, and $E[I_{10}]$ for four tie breaking rules: \mathcal{T}_{var} , \mathcal{T}_{maxs} , \mathcal{T}_{mins} , and \mathcal{T}_{meds} when $n=3$.

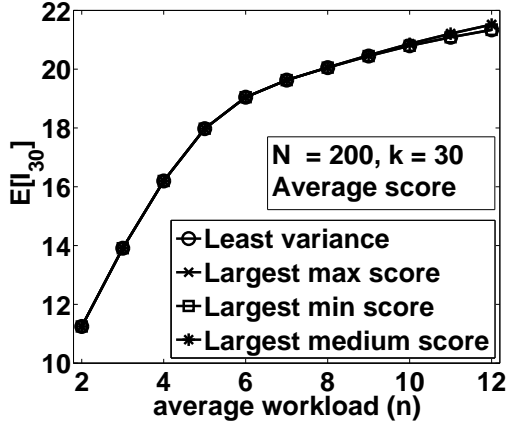
To evaluate the performance of these four tie breaking rules, let us select a voting rule: *average score rule* and we use expectation as our performance measure. We evaluate the performance of these four tie breaking rules on the conference recommendation system specified in Section 5.1. The numerical results of $E[I_{30}]$ are shown in Fig. 6. The numerical results of $E[I_1]$ and $E[I_5]$ for high *self-selectivity* papers are shown in Fig. 7. The numerical results of $E[I_1]$, $E[I_5]$, and $E[I_{10}]$ when $n=3$ are shown in Table 7.

In Fig. 6 and 7, the horizontal axis represents the number of reviews per paper, or n . The vertical axis shows the corresponding expectation. From Fig. 6, we could have the following observations. When we increase the reviewing workload, the expectation $E[I_{30}]$ corresponding to each tie breaking rule increased, which reflects that the accuracy of the conference recommendation system increased. From Fig. 6(a) and 6(c), we could see that when submitted papers are of high or low *self-selectivity*, the expectation curves corresponding to these four tie breaking rules overlapped together. In other words, these four rules have the same accuracy. From Fig. 6(b) and 6(d), we see that when submitted papers are of medium or random *self-selectivity*, the expectation curves corresponding to these four tie breaking rules bunched together and the *largest medium score rule* has a slightly higher accuracy than others. From Table 7, we see that when $n=3$, the chance of accepting the best paper is invariant of tie breaking rules, unless when submitted papers are of high *self-selectivity*, since values of $E[I_1]$ corresponding to medium, low or random *self-selectivity* are all around 0.98. This statement also holds for the top five and top ten papers. When submitted papers are of high *self-selectivity*, the chance of the best paper being accepted is low, or less than 0.6. The same statements holds for the top five and top ten papers. From Fig. 7, we see that for each tie breaking rule, we still have to increase the reviewing workload n to at least 7 so as to have a strong guarantee that the best paper or the top five papers will get accepted.

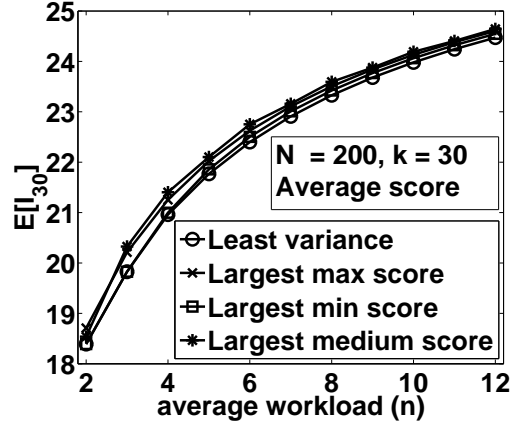
Lessons learned: These four tie breaking rules have comparable accuracy, no rule can outperform others remarkably. Thus, the improvement of conference recommendation system by tie breaking rules is limited. For each tie breaking rule, we till have to increase the reviewing workload n to at least seven to have a strong guarantee that the best paper will get in.

5.6 Effect of Reviewers Type – two types case

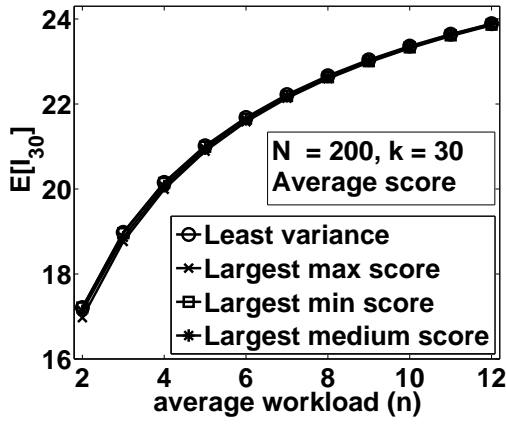
We extend the homogeneous conference recommendation system specified in Section 5.1 to a heterogeneous conference recommendation system, for which papers and reviewers can be of different types. Specifically, we consider papers and reviewers are of two types (e.g. system or theory). If paper-reviewer matches in the same type, the critical degree is high, else the critical degree is low. Here we explore the effect of the fraction that paper-reviewer matches in the same type on the accuracy of the recommendation. We use expectation as our performance measure and we set the number of review per paper in this heterogeneous recommendation system be four, or $n=4$. We vary the fraction of



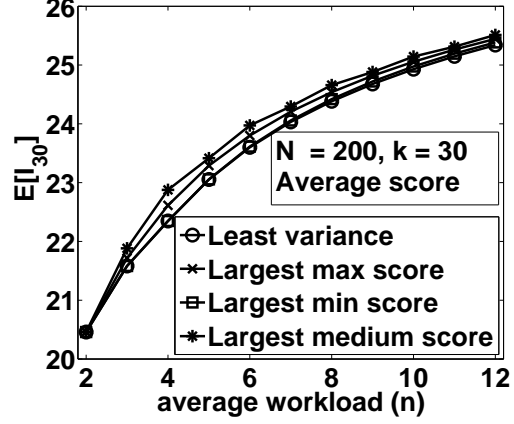
(a) High self-selectivity



(b) Medium self-selectivity

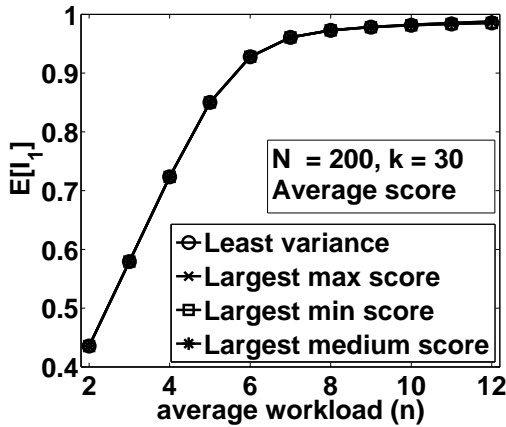


(c) Low self-selectivity

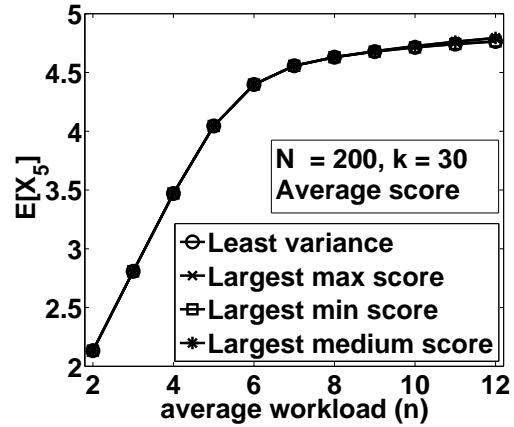


(d) Random self-selectivity

Figure 6: Expectation of $E[I_{30}]$ for four tie breaking rules: \mathcal{T}_{var} , \mathcal{T}_{maxs} , \mathcal{T}_{mins} , and \mathcal{T}_{meds} .



(a) # of top 1 paper get in



(b) # of top 5 papers get in

Figure 7: Expectation of $E[I_1]$, $E[I_5]$ for four tie breaking rules: \mathcal{T}_{var} , \mathcal{T}_{maxs} , \mathcal{T}_{mins} , and \mathcal{T}_{meds} for high self-selectivity submission.

paper-reviewer that matches in the same type from 0.1 to 1. Before showing our results, let us specify $\sigma(c_j^i)$ derived in Eq. (6): if papers-reviewer matches in the same type, $\sigma(c_j^i) = 0.5$, otherwise $\sigma(c_j^i) = 2$. The numerical results of $E[I_1]$, $E[I_5]$, $E[I_{10}]$, and $E[I_{30}]$ are shown in Fig. 8.

In Fig. 8, the horizontal axis represents the the fraction that paper-reviewer matches in the same type. The vertical axis shows the corresponding expectation. From Fig. 8, we have the following observations. When we increase the matching fraction, the expectation $E[I_1]$, $E[I_5]$, $E[I_{10}]$, and $E[I_{30}]$ increased. In other words, the accuracy of the conference recommendation system increased when more paper-reviewer matches in the same type. It is interesting to see that the expectation $E[I_1]$, $E[I_5]$, $E[I_{10}]$ and $E[I_{30}]$ increase in a linear rate and the rate corresponding to the high *self-selectivity* is the highest, and the random *self-selectivity* is the lowest. When papers are submitted with medium, low or random *self-selectivity*, the chance of accepting the best paper is invariant of the matching fraction, since the corresponding values of $E[I_1]$ are approximately equal to 1 when the matching fraction is 0 and increased slightly when matching fraction increased to 1. This statement also holds for the top five papers. When submitted papers are highly *self-selective*, the accuracy of the system is very sensitive to the matching fraction, this is because $E[I_1] \approx 0.5$, $E[I_{10}] \approx 5$ and $E[I_{30}] \approx 13$ when matching fraction is 0, and $E[I_1] \approx 0.9$, $E[I_{10}] \approx 8$ and $E[I_{30}] \approx 20$ when matching fraction is 1.

Lessons learned: If reviewers' preference and papers match, it can significantly affect the accuracy of the conference recommendation system. This is especially true for submitted papers which are of high *self-selectivity* (or prestigious conferences). To achieve this, it is especially important to select the appropriate technical program committee members to match the research topics since it can greatly influence the accuracy of the final recommendation.

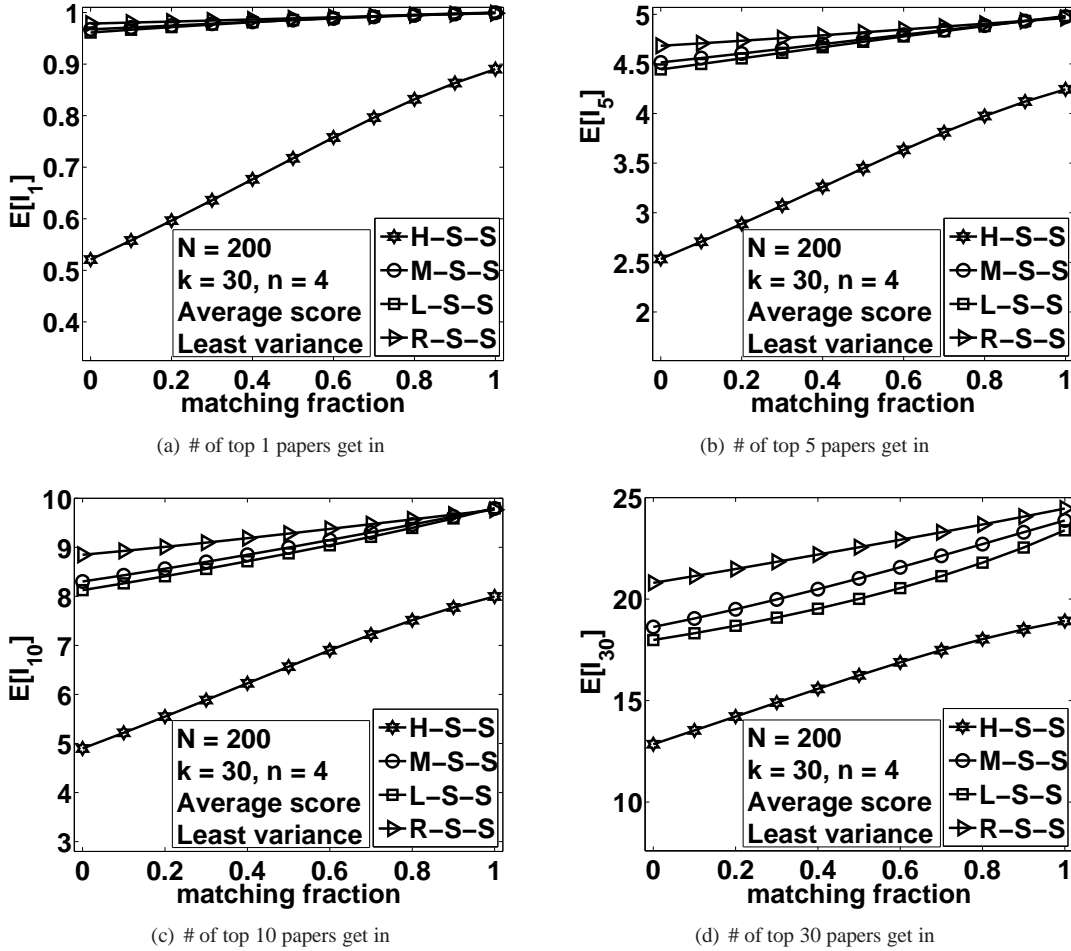


Figure 8: Impact of paper-reviewing matching

5.7 Effect of Reviewers type – many types case

We now generalize the types of papers or reviewers of a conference recommendation system specified in Section 5.6 to many types. We explore the effect of reviewer types on the accuracy of different voting rules. Specifically, we consider the following four representative voting rules: **Average score** (\mathcal{V}_{as}), **Eliminate the highest & lowest** (\mathcal{V}_{ehl}), **Punish low scores** (\mathcal{V}_{pl}) as given by Equation (27) – (29). We also consider the **Weighted average** (\mathcal{V}_{wa}) rule: The average score of paper P_i is weighted on e_j^i , where $j = 1, \dots, n_i$, the declared expertise level of reviewer $\mathcal{R}(S_j^i)$ who selects topics related to P_i , or

$$\gamma_i = \sum_{j=1}^{n_i} S_j^i e_j^i \left[\sum_{j=1}^{n_i} e_j^i \right]^{-1}. \quad (30)$$

We consider the least variance rule for tie breaking and we use expectation as our performance measure. In this evaluation, papers and reviewers are randomly matched. Before showing our results, let us specify two functions that related to model for reviewers types: the first one is $\sigma(c_j^i)$ within the probability distribution for score S_j^i derived in Eq. (6), which is specified by the following linear function

$$\sigma(c_j^i) = 0.5 + 1.5[1 - c_j^i],$$

the second one is a monotonic increasing function $f(\mu)$ within Eq. (9), which is specified by the following linear function

$$f(\mu) = \mu.$$

We set l to be 3, thus $e_j^i \in \{1, 2, 3\}$. The numerical results of $E[I_{30}]$ are shown Fig. 9.

In Fig. 9, the horizontal axis represents the number of reviews per paper, or n . The vertical axis shows the corresponding expectation. From Fig. 9, we have the following observations. When submitted papers are of high *self-selectivity*, the expectation curves corresponding to these four voting rules overlapped together. In other words, these four rules have the same accuracy for high *self-selectivity* submissions. When submitted papers of low *self-selectivity*, the *punish low scores rule* has the lowest accuracy than the other three voting rules which have nearly the same accuracy. When submitted papers of medium *self-selectivity*, the weighted average scoring rule and the *eliminate highest and lowest score rule* have nearly the same accuracy, and the weighted average scoring rule has slightly higher accuracy than the average scoring rule and *punish low scores rule*. This statement also holds for the submitted papers with random *self-selectivity*.

Lessons learned: When papers and reviewers are of many types and papers and reviewers are randomly matched, *weighted average score rule* can have slightly higher accuracy than *average score rule* and *punish low scores rule*, and it has nearly the same accuracy with *eliminate highest and lowest score rule*. We have the following similar observations in Section 5.4: these four voting rules have comparable accuracy, no rule can outperform others, thus the improvement of conference recommendation system by voting rules is limited. Again, there are number of interesting questions to explore, i.e., is the accuracy sensitive to anomaly behavior?

5.8 Effect of Anomaly Behavior – Random Scoring

Let us explore the effect of an anomaly behavior on the accuracy of a conference recommendation system. Specifically, we consider one potential anomaly behavior: *random scoring behavior*, under which a misbehaving reviewer gives a random score to any paper she reviews regardless of the quality of that paper. We vary the fraction of anomaly behavior from 0.1 to 1 and we use expectation as our performance measure. We explore the effect of anomaly behavior on the conference recommendation system specified in Section 5.1 with number of reviews per paper $n = 4$. The numerical results of $E[I_1]$, $E[I_5]$, $E[I_{10}]$, and $E[I_{30}]$ are shown in Fig. 10.

In Fig. 10, the horizontal axis represents the fraction of these misbehaving reviewers. The vertical axis shows the corresponding expectation. From Fig. 10, we have the following observations. When we increase the fraction of anomaly reviewers, the expectation decreased. In other words, the more anomaly reviewers, the lower is the accuracy of the conference recommendation system. It is interesting to note that the accuracy of the conference recommendation system decreases in a nearly linear rate. From Fig. 10(a), we see that for low *self-selectivity* papers, the chance of accepting the best paper can withstand a small fraction of misbehaving reviewers, say around 10%. This statement holds for the top five papers. When submitted papers are of high *self-selectivity*, around 40% of misbehaving reviewers can drastically disrupt the accuracy of a conference recommendation system. Because for this case, less than ten papers from the top 30 papers will get accepted, and less than four papers from the top ten papers will get accepted. Even

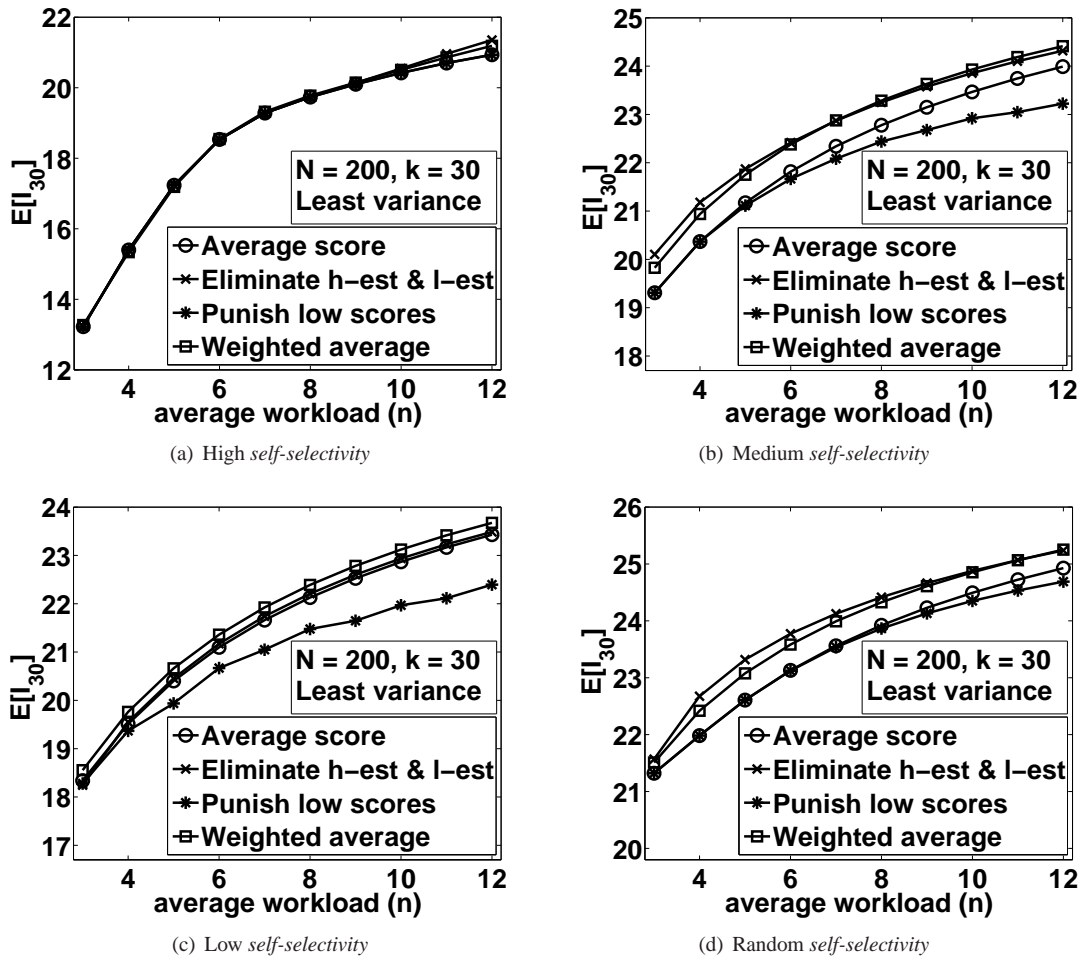


Figure 9: Expectation of I_{30} for four voting rules: \mathcal{V}_{as} , \mathcal{V}_{ehl} , \mathcal{V}_{pl} and \mathcal{V}_{wa} .

the best paper can only be accepted with the probability of less than 0.4. When papers are submitted with medium, low or random *self-selectivity*, around 60% misbehaving reviewers can drastically disrupt the accuracy of a conference recommendation system.

An interesting question is that which voting rule is more robust against this type of misbehaving reviewers? Here we evaluate three voting rules: *average score rule*, *eliminate highest and lowest* and *punish low scores* given by Equation (27) – (29). We use expectation as our performance measure. The numerical results of $E[I_{30}]$ are shown in Fig. 11.

In Fig. 11, the horizontal axis represents of the fraction of misbehaving reviewers. The vertical axis shows the corresponding expectation. From Fig. 11, we have the following observations. When we increase the fraction of anomaly reviewers, the expectation decreased. From Fig. 11(c), we see that when submitted papers are of low *self-selectivity*, the expectation curves corresponding to these three voting rules overlapped together. In other words, for the low *self-selectivity* papers, these three voting rules have the same robustness. From Fig. 11(b), 11(d), we see that when submitted papers are of medium or random *self-selectivity*, eliminate highest and lowest score rule are slightly more robust than the other two rules. From Fig. 11(a), we observe that when submitted papers of high *self-selectivity*, these three voting rules have the same robustness when the fraction of anomaly reviewer is less than 30%, and when it is higher than 30%, eliminate highest and lowest voting rule are more robust than the other two rules.

Lessons learned: Random anomaly behavior can significantly affect the accuracy of the conference recommendation system. This is especially true for submitted papers which are of high *self-selectivity* (or prestigious conferences). The conference recommendation system suffers significantly from this kind of anomaly behavior, say and 20% of misbehaving reviewers will reduce the probability of the best paper to be accepted to around 0.5. These four voting rules have comparable robustness, no rule can outperform others remarkably, thus to defend this kind of anomaly

behavior by voting rules may not be effective. Again, there are number of interesting questions to explore, i.e., how to improve the accuracy of the conference recommendation system?

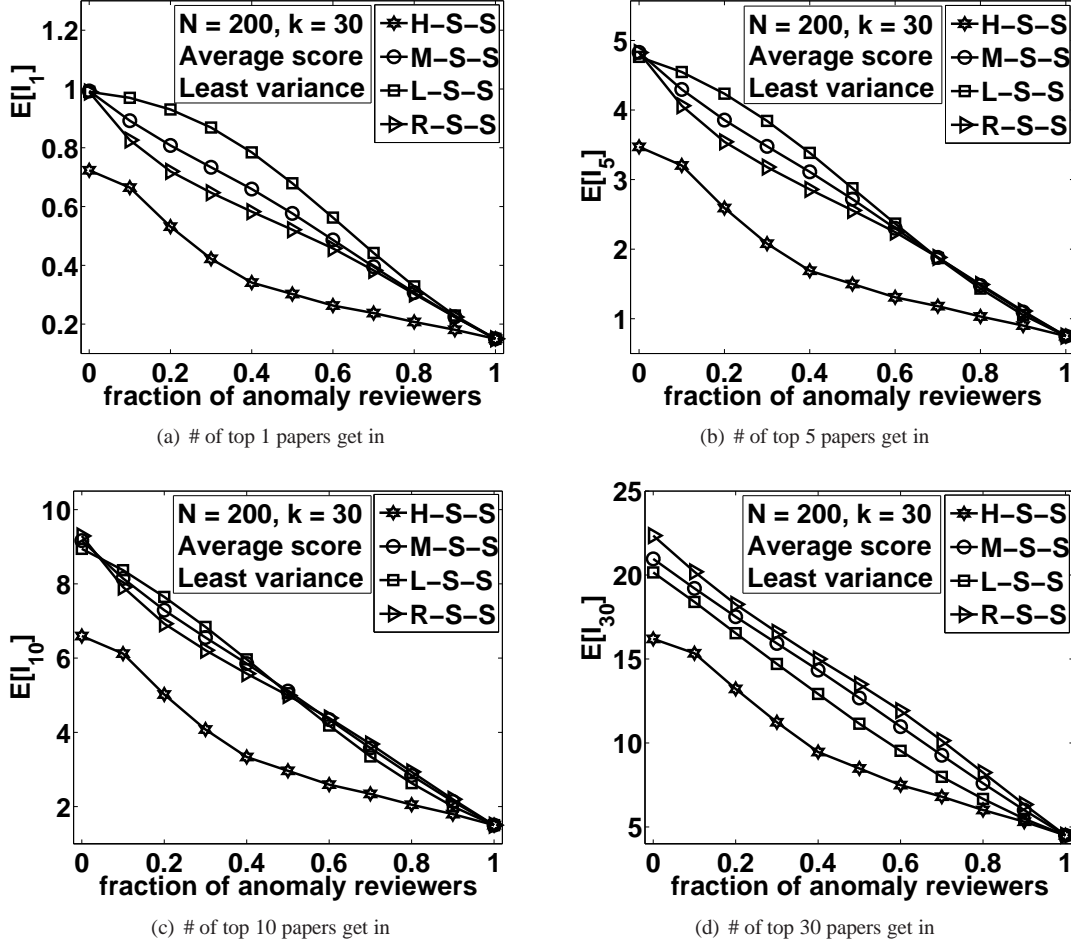


Figure 10: Impact of random anomaly behavior when $n = 4$.

5.9 Effect of Anomaly Behavior – Bias Scoring

Let us explore the effect of another potential anomaly behavior on the accuracy of a conference recommendation system. Specifically, we consider *bias scoring behavior*, under which a misbehaving reviewer gives a high score m to a paper if the evaluated quality is low, say less than three, otherwise gives a low score 1 to those papers whose evaluated quality is above three. We vary the fraction of anomaly behavior from 0 to 0.3 and we use expectation as our performance measure. We explore the effect of this type of anomaly behavior on the conference recommendation system specified in Section 5.1 with number of reviews per paper $n = 4$. The numerical results of $E[I_1]$, $E[I_5]$, $E[I_{10}]$, and $E[I_{30}]$ are shown in Fig. 12.

In Fig. 12, the horizontal axis represents the fraction of these misbehaving reviewers. The vertical axis shows the corresponding expectation. From Fig. 12, we have the following observations. When we increase the fraction of anomaly reviewers slightly, the expectation decreased remarkably. From Fig. 12(a), we see that for low *self-selectivity* papers, the chance of accepting the best paper can withstand a small fraction of misbehaving reviewers, say around 6%, but for the other three *self-selectivity* types, even a small fraction of misbehaving reviewers may lead to high inaccuracy. For example, when 6% of reviewers are misbehaving, the best paper only has less than 70% chance of being accepted for the *highly self-selective* paper submissions and it only has less than 80% chance of being accepted for the submitted papers which are of medium or random *self-selectivity*. Similar deterioration can be said for the top five papers. Around 15% of misbehaving reviewers can drastically disrupt the accuracy of a conference

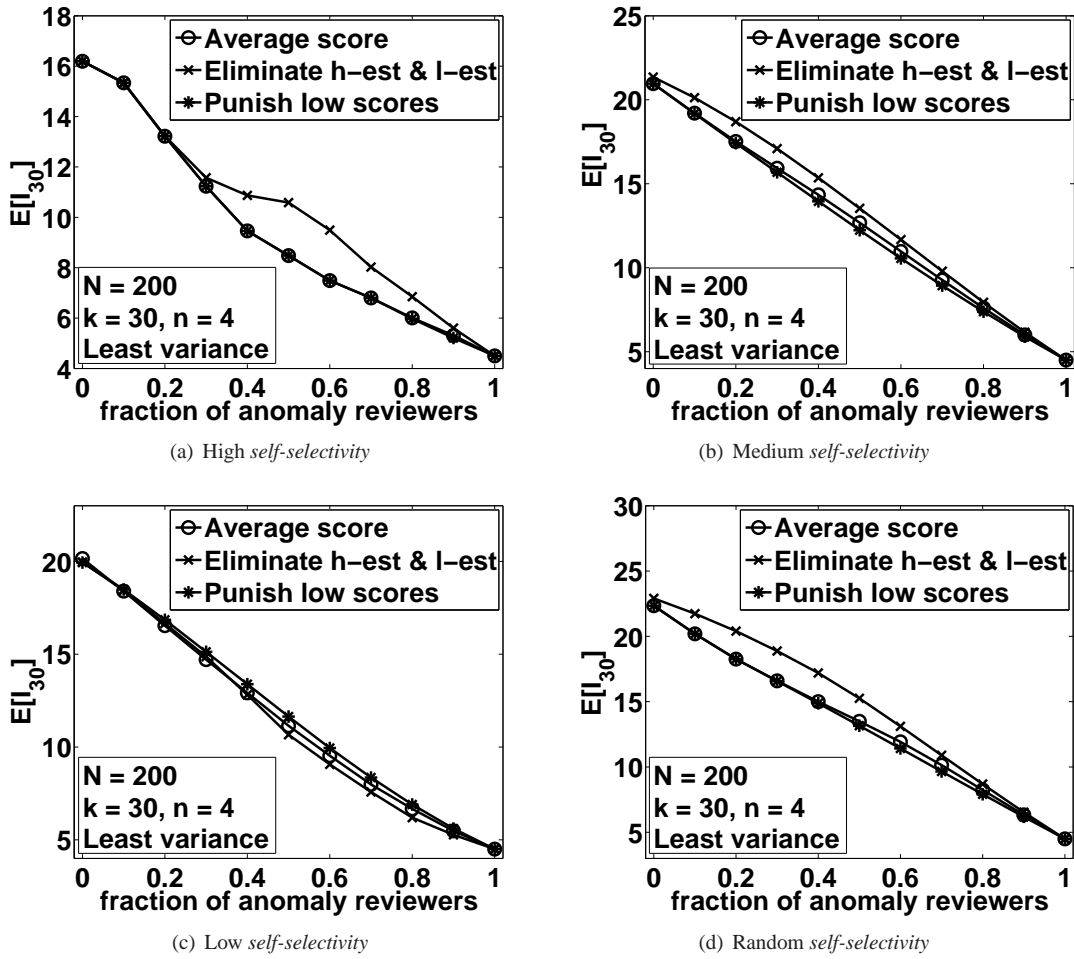


Figure 11: Robustness of three voting rules: \mathcal{V}_{as} , \mathcal{V}_{ehl} , and \mathcal{V}_{pl} when $n=4$.

recommendation system, since less than 15 papers from top 30 papers will get accepted.

Lessons learned: Bias scoring anomaly behavior can significantly affect the accuracy of the conference recommendation system. A small fraction of this kind of misbehaving reviewers can decrease the accuracy of the conference recommendation system dramatically. This is especially true for submitted papers which are of high, medium, or random *self-selectivity* (prestigious, medium, or newly started conferences). The conference recommendation system suffers severely from this kind of anomaly behavior, say 15% of misbehaving reviewers will disrupt the accuracy of the conference recommendation system.

5.10 Improving Conference Recommendation Systems

In reality, most conferences use homogeneous review strategy, with which each paper is reviewed by the *same number* of reviewers. One obvious advantage of this strategy is its fairness for all papers. But its efficiency in using the workload is low. Here we propose a *heterogeneous review strategy* that that can increase the efficiency of homogeneous review strategy. In other words, with the same reviewing workload W , it has a much higher chance to finally include the top k papers in the final acceptance recommendation. Assume that the the total reviewing workload for the homogeneous review strategy is $W = N * n$. Our heterogeneous review strategy works in two rounds:

Round 1: Eliminate half of the submitted papers using only half of the workload. Specifically, each paper will receive $\lfloor n/2 \rfloor$ reviews in the first round. After this reviewing round, apply a voting rule and tie breaking rule to eliminate $N/2$ papers.

Round 2: Select k papers from the remaining $N/2$ papers to accept. Each paper entering this round will receive $2\lfloor n/2 \rfloor$ reviews. After the reviewing process finished, combine their reviews in round 1 and round 2. Then apply a

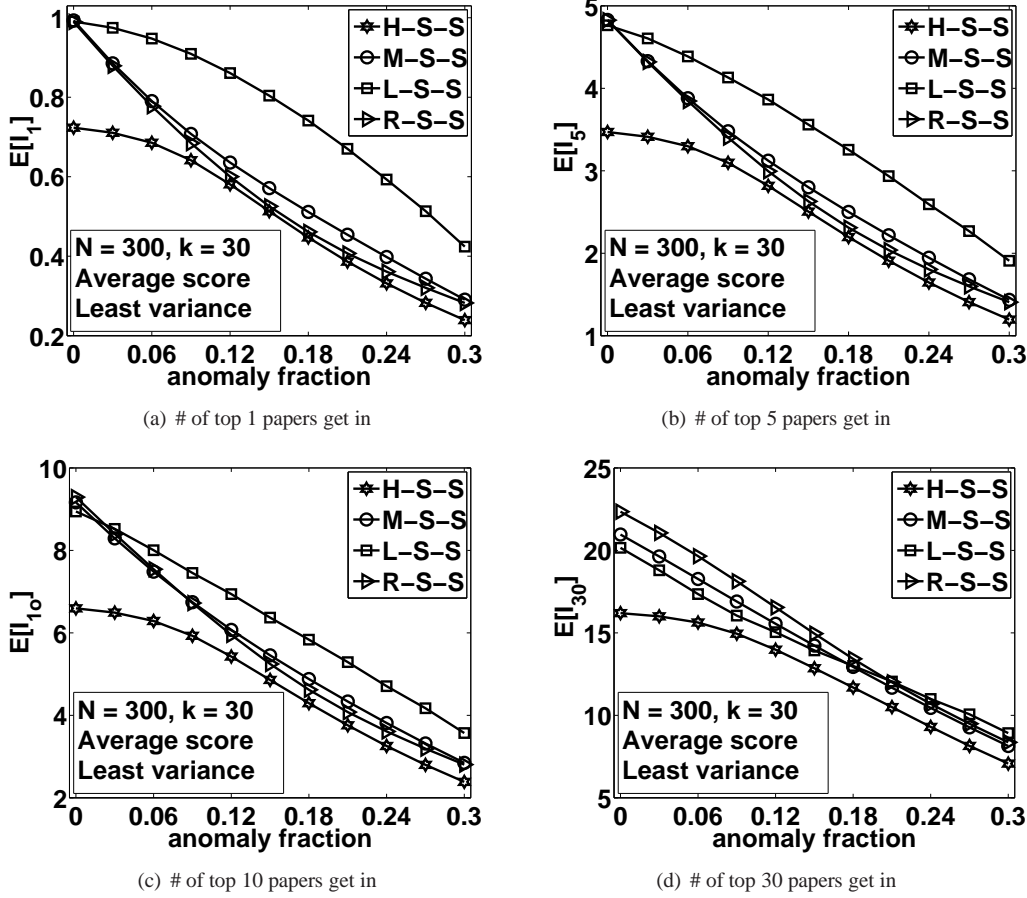


Figure 12: Impact of bias scoring anomaly behavior when $n = 4$.

voting rule and a tie breaking rule to select the top k papers to accept.

Definition 3 Let $E[I_i | \text{hom}]$ and $E[I_i | \text{hetero}]$ represent the expectation of I_i under homogeneous or heterogeneous review strategy applied respectively, where I_i is defined by Definition 2.

Definition 4 The improvement of heterogeneous review strategy over homogeneous review strategy is:

$$\Delta E[I_i] = E[I_i | \text{hetero}] - E[I_i | \text{hom}],$$

and the improvement ratio is:

$$\Delta E[I_i] / E[I_i | \text{hom}].$$

We evaluate these two strategies on a conference recommendation system specified in Section 5.1. The numerical results of $\Delta E[I_{30}]$ and $\Delta E[I_{30}] / E[I_{30} | \text{hom}]$ are shown in Fig. 13, where the horizontal axis represents the average reviewing workload n . The vertical axis shows the corresponding improvement or improvement ratio. From Fig. 13, we see an improvement of heterogeneous review strategy over homogeneous review strategy. When the reviewing workload is small, say $n = 3$, with heterogeneous review strategy at least one more paper from the top 30 papers will get accepted. For papers with high *self-selectivity*, we have more improvement wherein three or more papers from the top 30 papers will get accepted. When the average reviewing workload increased to six, the improvement becomes stabilized. When the reviewing workload is three, the improvement is the highest, or around four more papers from the top 30 papers will get accepted for papers submitted with high *self-selectivity*, and the improvement ratio for this case is around 30%.

An interesting question is that with heterogeneous review strategy, how large the average reviewing workload do we need. We apply our heterogeneous strategy to the conference recommendation system specified in Section 5.3. We

use expectation as performance measure. The numerical results of $E[I_1]$, $E[I_5]$, $E[I_{10}]$ and $E[I_{30}]$ are shown in Fig. 14, where the horizontal axis represents the average reviewing workload n . The vertical axis shows the corresponding expectation. From Fig. 14, we see that we need to increase the workload to at least five such that we have a strong guarantee that the best paper will get accepted, which *reduces* the average review workload by two as compared with the homogeneous review strategy, in which we need to increase the average reviewing workload to at least seven as stated in Section 5.3.

Summary: Lessons learned: Our heterogeneous review strategy uses equal or less resource (e.g., reviewing workload) than the homogeneous review strategy and at the same time, achieve higher accuracy.

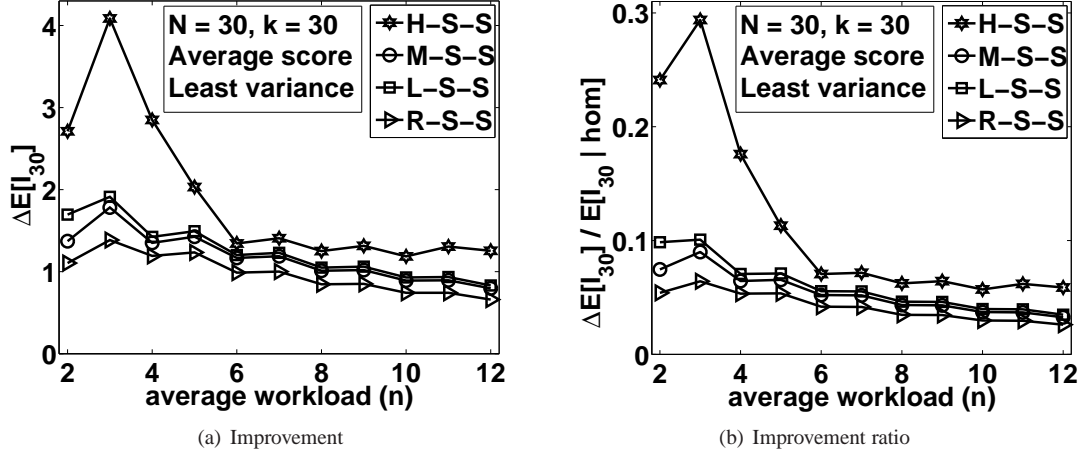


Figure 13: Improvement of heterogeneous review strategy on homogeneous review strategy

6 Related Work

In [8, 9, 14, 21], authors studied peer review systems. Typically, the main issue is the reviewer assignment problem which contains three phrases: specifying the assignment constraint, computing the matching degree between reviewers and submissions, and optimizing the assignment with constraints. Disciplines like information retrieval [8], artificial intelligence [11, 21] and operations research [5, 9, 14], etc address these assignment problems.

Authors in [4, 19, 24, 25] worked on the *group recommendation systems* and address issues on rating scale and [2, 15, 6] on preference aggregation. Rating is used to show individuals' preferences, and in [25], authors stated that discrete rating scales (number of rating points) outperform continuous rating scales. In [4], authors evaluated the reliability of rating scales and showed evidence that more rating points will have a more reliable rating. In [19], authors stated that the best rating scale is around five to ten rating points. Preference aggregation is the process to merge the preference of multiple people so as to make recommendations. Basically the aggregation method can be divided into two classes based on the preference type: cardinal ranking or ordinal ranking. For cardinal ranking case, *weighted average strategy*[16] is the most popular strategy, and it is used in *PolyLen*. The second class is the ordinal ranking preference, for which, each individual's preference is shown by a ranked list of a subset of the candidates. For this case, users' preferences are treated as a set of constraints and a preference aggregation approach attempts to find recommendations that satisfy the constraints of all users[2, 6, 15]. As far as we know, our work is the first that study the mathematical modeling of competitive group recommendation systems and apply it to peer review systems.

7 Conclusions

This is the first paper that provides a mathematical model and analysis on a competitive group recommendation system. We apply it to a conference peer review system and show how various factors may influence the overall accuracy of the final recommendation. We formally analyze the model and through this analysis, we gain the insight on developing a randomized algorithm, which is both computationally efficient and can provide performance guarantees on various performance measures. Number of interesting observations are found, e.g., for a medium tier conference, three reviews

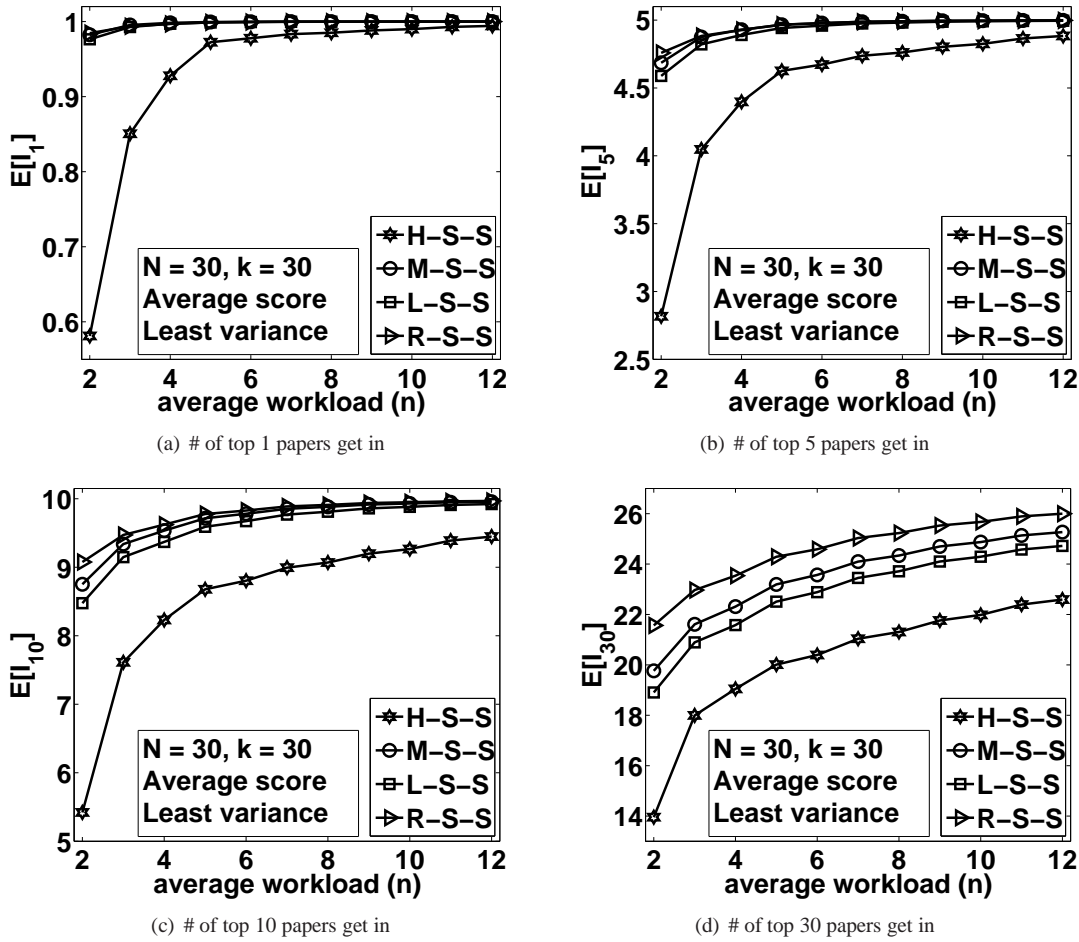


Figure 14: Expectation of I_1 , I_5 , I_{10} and I_{30} with heterogeneous review strategy.

per paper are sufficient to achieve high accuracy in the final recommendation, but for some prestigious conferences, we need at least seven reviews per paper. Lastly, we propose a heterogeneous review strategy that requires equal or less reviewing workload but can produce a more accurate recommendation than the homogeneous review strategy. We believe our model and methodology are important building blocks for researchers to study competitive group recommendation systems.

References

- [1] S. Amer-Yahia, S. Roy, A. Chawlat, G. Das, and C. Yu. Group recommendation: Semantics and efficiency. *Proc. of VLDB*, 2009.
- [2] J. Baskin and S. Krishnamurthi. Preference aggregation in group recommender systems for committee decision-making. In *Proc. of ACM RecSys*, 2009.
- [3] L. Boratto and S. Carta. State-of-the-art in group recommendation and new approaches for automatic identification of groups. *Information Retrieval and Mining in Distributed Environments*, pages 1–20, 2011.
- [4] G. Churchill Jr and J. Peter. Research design effects on the reliability of rating scales: a meta-analysis. *Journal of Marketing Research*, pages 360–375, 1984.
- [5] W. Cook, B. Golany, M. Kress, M. Penn, and T. Raviv. Optimal allocation of proposals to reviewers to facilitate effective ranking. *Management Science*, pages 655–661, 2005.

- [6] W. Cook, B. Golany, M. Penn, and T. Raviv. Creating a consensus ranking of proposals from reviewers partial ordinal rankings. *Computers & operations research*, 34(4):954–965, 2007.
- [7] M. Deshpande and G. Karypis. Item-based top-n recommendation algorithms. *ACM TOIS*, 22(1):143–177, 2004.
- [8] S. Dumais and J. Nielsen. Automating the assignment of submitted manuscripts to reviewers. In *Proc. of ACM SIGIR*, 1992.
- [9] D. Hartvigsen, J. Wei, and R. Czuchlewski. The conference paper-reviewer assignment problem. *Decision Sciences*, 30(3):865–876, 1999.
- [10] J. Herlocker, J. Konstan, L. Terveen, and J. Riedl. Evaluating collaborative filtering recommender systems. *ACM TOIS*, 22(1):5–53, 2004.
- [11] S. Hettich and M. Pazzani. Mining for proposal reviewers: lessons learned at the national science foundation. In *Proc. of SIGKDD’06*.
- [12] A. Jameson. More than the sum of its members: challenges for group recommender systems. In *Proc. of ACM working conference on Advanced visual interfaces*, 2004.
- [13] A. Jameson, S. Baldes, and T. Kleinbauer. Two methods for enhancing mutual awareness in a group recommender system. In *Proceedings of ACM working conference on Advanced visual interfaces*, 2004.
- [14] M. Karimzadehgan and C. Zhai. Constrained multi-aspect expertise matching for committee review assignment. In *Proc. of CIKM’09*.
- [15] F. Lorenzi, F. dos Santos, P. Ferreira, and A. Bazzan. Optimizing preferences within groups: a case study on travel recommendation. *Advances in Artificial Intelligence-SBIA*, pages 103–112, 2008.
- [16] J. Masthoff. Group modeling: Selecting a sequence of television items to suit a group of viewers. *User Modeling and User-Adapted Interaction*, 14(1):37–85, 2004.
- [17] M. Mitzenmacher and E. Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge Univ Pr, 2005.
- [18] M. Oconnor, D. Cosley, J. Konstan, and J. Riedl. Polylens: A recommender system for groups of users. In *Proc. of ECSCW 2001*.
- [19] C. Preston and A. Colman. Optimal number of response categories in rating scales: reliability, validity, discriminating power, and respondent preferences. *Acta psychologica*, 104(1):1–15, 2000.
- [20] P. Resnick and H. Varian. Recommender systems. *Communications of the ACM*, 40(3):56–58, 1997.
- [21] M. Rodriguez and J. Bollen. An algorithm to determine peer-reviewers. In *Proc. of ACM CIKM*, 2008.
- [22] B. Sarwar, G. Karypis, J. Konstan, and J. Reidl. Item-based collaborative filtering recommendation algorithms. In *Proc. of WWW’01*.
- [23] J. Schafer, J. Konstan, and J. Riedi. Recommender systems in e-commerce. In *Proc. of ACM EC*, 1999.
- [24] E. I. Sparling and S. Sen. Rating: how difficult is it? In *Proc. of ACM RecSys*, 2011.
- [25] E. Svensson. Comparison of the quality of assessments using continuous and discrete ordinal rating scales. *Biometrical Journal*, 42(4):417–434, 2000.

Appendix

Theorem 7 (Chernoff Bound[17]) *Let X_1, \dots, X_n be independent random variables with $X_i = 1$ with probability p and 0 otherwise. Let $X = \sum_{i=1}^n X_i$ and let $\mu = E[X] = np$. Then we have*

$$\Pr[|X - \mu| \geq \epsilon\mu] \leq 2e^{-\mu\epsilon^2/3}.$$

In the following lemma we derive a loose bound on the number of simulation rounds K but have a good performance guarantee on the pmd of $I(k)$, or $\Pr[I(k) = i]$, $\forall i$ by carefully applying Theorem 7.

Lemma 3 *When the following condition holds:*

$$K \geq \max_{i \in \{0, \dots, k\}} \frac{3 \ln(2(k+1)/\delta)}{\Pr[I(k) = i]\epsilon^2}, \quad (31)$$

then for each $i = 0, 1, \dots, k$,

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \Pr[I(k) = i],$$

holds with probability at least $1 - \delta/(k+1)$.

Proof: Without lose of any generality, consider the performance guarantee on the approximation of $\Pr[I(k) = i]$, where $i = 0, 1, \dots, k$. Our goal is to show

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \Pr[I(k) = i].$$

Let \mathbf{I}_{ij} be an indicator random variable defined by

$$\mathbf{I}_{ij} = \begin{cases} 1 & \text{if in } j\text{-th round, } |\mathcal{A}^I(k) \cap \mathcal{A}(k)| = i \\ 0 & \text{otherwise} \end{cases},$$

where $j = 1, \dots, K$. There are two cases to explore:

Case 1: $\Pr[I(k) = i] = 0$. The physical meaning implies that the event $I(k) = i$ never happen, which result in $\mathbf{I}_{ij} = 0$ for all $j = 0, 1, \dots, K$. Hence, $\widehat{\Pr}[I(k) = i] = \sum_{j=1}^K \mathbf{I}_{ij} / K = 0$. Then we have

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \Pr[I(k) = i].$$

Case 2: $\Pr[I(k) = i] \neq 0$. Since each round runs independently, thus random variables $\mathbf{I}_{i1}, \dots, \mathbf{I}_{iK}$ are independent random variables with $\mathbf{I}_{ij} = 1$ with probability $\Pr[I(k) = i]$ and 0 otherwise. Let $\mathbf{I}_i = \sum_{j=1}^K \mathbf{I}_{ij}$. Then $E[\mathbf{I}_i] = K\Pr[I(k) = i]$. From Algorithm 1 we could see that $\widehat{\Pr}[I(k) = i]$, the approximate value of $\Pr[I(k) = i]$, is given by $\widehat{\Pr}[I(k) = i] = \mathbf{I}_i / K$. Then by applying Theorem 7 we have,

$$\begin{aligned} & \Pr \left[\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \geq \epsilon \Pr[I(k) = i] \right] \\ &= \Pr \left[\left| \mathbf{I}_i - K\Pr[I(k) = i] \right| \geq \epsilon K \Pr[I(k) = i] \right] \\ &\leq 2e^{-K\Pr[I(k) = i]\epsilon^2/3}, \end{aligned}$$

by substituting K with Inequality (31), we have

$$\begin{aligned} & \Pr \left[\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \geq \epsilon \Pr[I(k) = i] \right] \\ &\leq \delta/(k+1). \end{aligned}$$

Finally the proof of this lemma can be completed by

$$\begin{aligned} & \Pr \left[\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \Pr[I(k) = i] \right] \\ & \geq 1 - \Pr \left[\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \geq \epsilon \Pr[I(k) = i] \right] \\ & \geq 1 - \delta / (k + 1). \end{aligned}$$

This lemma is proved. ■

Lemma 3 shows the performance guarantee on $\Pr[I(k) = i]$ with success probability at least $1 - \delta / (k + 1)$ for each specific $i = 0, 1, \dots, k$. Then, what is the success probability for all $i = 0, 1, \dots, k$? The answer of this question is stated in the following lemma.

Lemma 4 *When the following condition holds:*

$$K \geq \max_{\substack{i \in \{0, \dots, k\} \\ \Pr[I(k) = i] \neq 0}} \frac{3 \ln(2(k + 1)/\delta)}{\Pr[I(k) = i] \epsilon^2},$$

then

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \Pr[I(k) = i],$$

holds for all $x = 0, 1, \dots, k$, with probability at least $1 - \delta$.

Proof: Let E_i denote the event that $\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \Pr[I(k) = i]$ holds. From Lemma 3 we see that for each $i = 0, 1, \dots, k$, the condition

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \Pr[I(k) = i]$$

holds with probability at least $1 - \delta / (k + 1)$, thus

$$\Pr[E_i] \geq 1 - \delta / (k + 1).$$

Our goal is to derive the probability that condition

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \Pr[I(k) = i]$$

holds for all $i = 0, 1, \dots, k$. Specifically, the probability that events E_0, E_1, \dots, E_k all happens, or

$$\Pr[E_0, \dots, E_k].$$

Note that the physical meaning of E_i is that $\widehat{\Pr}[I(k) = i]$, the approximate value of $\Pr[I(k) = i]$, is close to the true value of $\Pr[I(k) = i]$. Thus the physical meaning of $\cap_{i \in \mathcal{F}} E_i$, where $\mathcal{F} \subseteq \{0, \dots, k\}$, is that the approximate value of $\Pr[I(k) = i]$, is close to the true value of $\Pr[I(k) = i]$, for all $i \in \mathcal{F}$. Thus $\cap_{i \in \mathcal{F}} E_i$ happens implies that $\sum_{i \in \mathcal{F}} \widehat{\Pr}[I(k) = i]$, the approximate value of $\sum_{i \in \mathcal{F}} \Pr[I(k) = i]$ is close to the real value of $\sum_{i \in \mathcal{F}} \Pr[I(k) = i]$. Since

$$\sum_{i \in \overline{\mathcal{F}}} \widehat{\Pr}[I(k) = i] - \sum_{i \in \overline{\mathcal{F}}} \Pr[I(k) = i] = \sum_{i \in \mathcal{F}} \widehat{\Pr}[I(k) = i] - \sum_{i \in \mathcal{F}} \Pr[I(k) = i],$$

where $\overline{\mathcal{F}} = \{0, 1, \dots, k\} \setminus \mathcal{F}$ is the complement of \mathcal{F} , thus we have that $\sum_{i \in \mathcal{F}} \widehat{\Pr}[I(k) = i]$ is close to the real value of $\sum_{i \in \mathcal{F}} \Pr[I(k) = i]$ if and only if $\sum_{i \in \overline{\mathcal{F}}} \widehat{\Pr}[I(k) = i]$ is close the real value of $\sum_{i \in \overline{\mathcal{F}}} \Pr[I(k) = i]$. Thus $\cap_{i \in \mathcal{F}} E_i$ implies that $\sum_{i \in \overline{\mathcal{F}}} \widehat{\Pr}[I(k) = i]$ is close the real value of $\sum_{i \in \overline{\mathcal{F}}} \Pr[I(k) = i]$.

The event $\cap_{i \in \mathcal{G}} E_i$, where $\mathcal{G} \subseteq \overline{\mathcal{F}}$, is more likely to happen given the prior information that $\sum_{i \in \overline{\mathcal{F}}} \widehat{\Pr}[I(k) = i]$ is close the real value of $\sum_{i \in \overline{\mathcal{F}}} \Pr[I(k) = i]$ than given noting at all. Since $\cap_{i \in \mathcal{F}} E_i$ implies that $\sum_{i \in \overline{\mathcal{F}}} \widehat{\Pr}[I(k) = i]$ is close the real value of $\sum_{i \in \overline{\mathcal{F}}} \Pr[I(k) = i]$, thus $\cap_{i \in \mathcal{G}} E_i$, where $\mathcal{G} \subseteq \overline{\mathcal{F}}$, is more likely to happen given prior information that $\cap_{i \in \mathcal{F}} E_i$ happens than given nothing at all. Or mathematically,

$$\Pr[\cap_{i \in \mathcal{G}} E_i \mid \cap_{i \in \mathcal{F}} E_i] \geq \Pr[\cap_{i \in \mathcal{G}} E_i],$$

where $\mathcal{G} \subseteq \overline{\mathcal{F}}$. Based on this fact, we have

$$\Pr[E_0, \dots, E_k] = \prod_{i=0}^k \Pr[E_i | \cap_{j \in \mathcal{F}_i} E_j] \geq \prod_{i=0}^k \Pr[E_i],$$

where $\mathcal{F}_i = \{0, \dots, k\} \setminus \{0, \dots, i\}$. By substituting $\Pr[E_i]$ with $\Pr[E_i] \geq 1 - \delta/(k+1)$ we have

$$\Pr[E_0, \dots, E_k] \geq [1 - \delta/(k+1)]^{k+1} \geq 1 - \delta,$$

which completes the proof. \blacksquare

In the following lemma we derive the error bound of expectation, or $|\widehat{E}[I(k)] - E[I(k)]|$, under the condition that $|\widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]| \leq \epsilon \Pr[I(k) = i]$ holds for all $i = 0, 1, \dots, k$.

Lemma 5 *When the following condition holds:*

$$|\widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]| \leq \epsilon \Pr[I(k) = i],$$

for all $i = 0, 1, \dots, k$, then

$$|\widehat{E}[I(k)] - E[I(k)]| \leq \epsilon E[I(k)]$$

holds.

Proof: The proof is quite straightforward,

$$\begin{aligned} |\widehat{E}[I(k)] - E[I(k)]| &= \left| \sum_{i=0}^k i (\widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]) \right| \\ &\leq \sum_{i=0}^k i |\widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]| \\ &\leq \sum_{i=0}^k \epsilon i \Pr[I(k) = i] = \epsilon E[I(k)], \end{aligned}$$

which completes the proof. \blacksquare

In the following, we let $\Delta E[I(k)] = \widehat{E}[I(k)] - E[I(k)]$ and $\Delta p_i = \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]$. In the following we first derive the bound of $(\Delta E[I(k)])^2$, and then apply the bound of $(\Delta E[I(k)])^2$ to derive the bound of variance, or $|\widehat{\text{Var}}[I(k)] - \text{Var}[I(k)]|$. The bound of $(\Delta E[I(k)])^2$ is stated in the following lemma.

Lemma 6 *When the following condition holds:*

$$|\widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]| \leq \epsilon \Pr[I(k) = i]$$

for all $i = 0, 1, \dots, k$, then

$$(\Delta E[I(k)])^2 \leq \epsilon^2 \text{Var}[I(k)] \quad (32)$$

holds.

Proof: First we have $|\Delta p_i| \leq \epsilon \Pr[I(k) = i]$. It is straightforward to show $\sum_{i=0}^k \Delta p_i = 0$, or

$$\begin{aligned} \sum_{i=0}^k \Delta p_i &= \sum_{i=0}^k \widehat{\Pr}[I(k) = i] - \sum_{i=0}^k \Pr[I(k) = i] \\ &= 0. \end{aligned}$$

Based on this fact, we could have

$$\begin{aligned} (\Delta E[I(k)])^2 &= \left(\sum_{i=0}^k i \Delta p_i \right)^2 \\ &= \left(\sum_{i=0}^k i \Delta p_i - E[I(k)] \sum_{i=0}^k \Delta p_i \right)^2 \\ &\leq \left(\sum_{i=0}^k |i - E[I(k)]| |\Delta p_i| \right)^2. \end{aligned} \quad (33)$$

Then by substituting $|\Delta p_i|$ with $|\Delta p_i| \leq \epsilon \Pr[I(k) = i]$ we have

$$(\Delta E[I(k)])^2 \leq \left(\sum_{i=0}^k |i - E[X]| \epsilon \Pr[I(k) = i] \right)^2$$

by applying Cauchy's Inequality we have,

$$\begin{aligned} (\Delta E[I(k)])^2 &\leq \left(\sum_{i=0}^k \epsilon^2 \Pr[I(k) = i] \right) \left(\sum_{i=0}^k (i - E[X])^2 \Pr[I(k) = i] \right) \\ &= \epsilon^2 \text{Var}[I(k)], \end{aligned}$$

which completes the proof. ■

In the following lemma we apply Lemma 6 to derive the error bound of variance, or $|\widehat{\text{Var}}[I(k)] - \text{Var}[I(k)]|$.

Lemma 7 *When the following condition holds:*

$$\left| \widehat{Pr}[I(k) = i] - Pr[I(k) = i] \right| \leq \epsilon \Pr[I(k) = i]$$

for all $i = 0, 1, \dots, k$, then

$$\left| \widehat{\text{Var}}[I(k)] - \text{Var}[I(k)] \right| \leq \epsilon(1 + \epsilon) \text{Var}[I(k)]$$

holds.

Proof: First we can write $\widehat{\text{Var}}[I(k)] - \text{Var}[I(k)]$ in the following form:

$$\begin{aligned} \widehat{\text{Var}}[I(k)] - \text{Var}[I(k)] &= \sum_{i=0}^k i^2 \Delta p_i - 2E[I(k)] \Delta E[I(k)] - (\Delta E[I(k)])^2 \\ &= \sum_{i=0}^k (i - E[I(k)])^2 \Delta p_i - (E[I(k)])^2 \sum_{i=0}^k \Delta p_i - (\Delta E[I(k)])^2 \end{aligned}$$

Since $\sum_{i=0}^k \Delta p_i = 0$, thus we have

$$\begin{aligned} \left| \widehat{\text{Var}}[I(k)] - \text{Var}[I(k)] \right| &= \left| \sum_{i=0}^k (i - E[I(k)])^2 \Delta p_i - (\Delta E[I(k)])^2 \right| \\ &\leq \sum_{i=0}^k (i - E[I(k)])^2 |\Delta p_i| + (\Delta E[I(k)])^2, \end{aligned} \tag{34}$$

by substituting Δp_i with $|\Delta p_i| \leq \epsilon \Pr[I(k) = i]$, and substituting $(\Delta E[I(k)])^2$ with Inequality (32) we have

$$|\widehat{\text{Var}}[I(k)] - \text{Var}[I(k)]| \leq \epsilon(1 + \epsilon) \text{Var}[I(k)],$$

which completes the proof. ■

In the following lemmas, we show how to tradeoff the number of simulation rounds K with the performance guarantees. The tradeoff between K and the performance guarantee on pmf is stated in the following theorem.

Lemma 8 *When the following condition holds:*

$$K \geq 3 \ln(2(k+1)/\delta) / \epsilon^2, \tag{35}$$

then for each $i = 0, \dots, k$,

$$\left| \widehat{Pr}[I(k) = i] - Pr[I(k) = i] \right| \leq \epsilon \sqrt{Pr[I(k) = i]},$$

holds with probability at least $1 - \delta/(k+1)$.

Proof: Without lose of any generality, consider the performance guarantee on the approximation of $\Pr[I(k) = i]$, where $i = 0, 1, \dots, k$. Our goal is to show

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \sqrt{\Pr[I(k) = i]}.$$

Following similar method in Lemma 3 and the same notations defined in the proof of Lemma 3, we could have the following:

Case 1: $\Pr[I(k) = i] = 0$. Then $\widehat{\Pr}[I(k) = i] = 0$. Hence

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \sqrt{\Pr[I(k) = i]}.$$

Case 2: $\Pr[I(k) = i] \neq 0$. Following similar method in Lemma 3 we have

$$\begin{aligned} & \Pr \left[\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \sqrt{\Pr[I(k) = i]} \right] \\ & \geq 1 - \Pr \left[\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \geq \epsilon \sqrt{\Pr[I(k) = i]} \right] \\ & = 1 - \Pr \left[\left| \mathbf{I}_i - K \Pr[I(k) = i] \right| \geq \frac{\epsilon K \Pr[I(k) = i]}{\sqrt{\Pr[I(k) = i]}} \right] \\ & \geq 1 - 2e^{-K\epsilon^2/3}, \end{aligned}$$

by substituting K with Inequality (35) we can prove this lemma. ■

Using the same method in Lemma 3, we can prove the following lemma.

Lemma 9 When $K \geq 3 \ln(2(k+1)/\delta)/\epsilon^2$, then

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \sqrt{\Pr[I(k) = i]}$$

holds for all $i = 0, 1, \dots, k$ with probability at least $1 - \delta$.

In the following lemma we derive the error bound of expectation, or $|\widehat{E}[I(k)] - E[I(k)]|$, under the condition that $|\widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]| \leq \epsilon \sqrt{\Pr[I(k) = i]}$ holds for all $i = 0, 1, \dots, k$.

Lemma 10 When the following condition holds:

$$\left| \widehat{\Pr}[I(k) = i] - \Pr[I(k) = i] \right| \leq \epsilon \sqrt{\Pr[I(k) = i]}$$

holds for all $i = 0, 1, \dots, k$, then

$$\left| \widehat{E}[I(k)] - E[I(k)] \right| \leq \epsilon \sqrt{(k+1)kE[I(k)]/2}$$

holds.

Proof: First we have $\Delta p_i \leq \epsilon \sqrt{\Pr[I(k) = i]}$. Then we have the following

$$\begin{aligned} \left| \widehat{E}[I(k)] - E[I(k)] \right| & \leq \sum_{i=0}^k i |\Delta p_i| \\ & \leq \sum_{i=0}^k \epsilon i \sqrt{\Pr[I(k) = i]}, \end{aligned}$$

by applying Cauchy's Inequality we have,

$$\begin{aligned} \left| \widehat{E}[I(k)] - E[I(k)] \right| & \leq \epsilon \sqrt{\left(\sum_{i=0}^k i \right) \left(\sum_{i=0}^k i \Pr[I(k) = i] \right)} \\ & = \epsilon \sqrt{(k+1)kE[I(k)]/2}, \end{aligned}$$

which completes the proof. ■

In the following lemma we derive the bound of $(\Delta E[I(k)])^2$ under the condition that $|\widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]| \leq \epsilon \sqrt{\Pr[I(k) = i]}$ holds for all $i = 0, 1, \dots, k$.

Lemma 11 *When the following condition holds:*

$$|\widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]| \leq \epsilon \sqrt{\Pr[I(k) = i]}$$

holds for all $i = 0, 1, \dots, k$, then

$$(\Delta E[I(k)])^2 \leq \epsilon^2(k+1)\text{Var}[I(k)] \quad (36)$$

holds.

Proof: First we have $\Delta p_i \leq \epsilon \sqrt{\Pr[I(k) = i]}$. From Inequality (33) we have

$$(\Delta E[I(k)])^2 \leq \left(\sum_{i=0}^k |i - E[I(k)]| |\Delta p_i| \right)^2,$$

by substituting $|\Delta p_i|$ with $\Delta p_i \leq \epsilon \sqrt{\Pr[I(k) = i]}$ we have

$$(\Delta E[I(k)])^2 \leq \left(\sum_{i=0}^k |i - E[I(k)]| \epsilon \sqrt{\Pr[I(k) = i]} \right)^2,$$

by applying Cauchy's Inequality we have,

$$\begin{aligned} (\Delta E[I(k)])^2 &\leq \left(\sum_{i=0}^k \epsilon^2 \right) \left(\sum_{i=0}^k (i - E[I(k)])^2 \Pr[I(k) = i] \right) \\ &= \epsilon^2(k+1)\text{Var}[I(k)], \end{aligned}$$

which completes the proof. ■

In the following lemma we apply Lemma 11 to derive the error bound of variance, or $|\widehat{\text{Var}}[I(k)] - \text{Var}[I(k)]|$, under the condition that $|\widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]| \leq \epsilon \sqrt{\Pr[I(k) = i]}$ holds for all $i = 0, 1, \dots, k$.

Lemma 12 *When the following condition holds:*

$$|\widehat{\Pr}[I(k) = i] - \Pr[I(k) = i]| \leq \epsilon \sqrt{\Pr[I(k) = i]}$$

holds for all $i = 0, 1, \dots, k$, then

$$|\widehat{\text{Var}}[I(k)] - \text{Var}[I(k)]| \leq \epsilon(k+1) \left(\sqrt{(2k+1)\text{Var}[I(k)]/6} + \epsilon \text{Var}[I(k)] \right)$$

holds.

Proof: First we have $\Delta p_i \leq \epsilon \sqrt{\Pr[I(k) = i]}$. From Inequality (34) we have

$$|\widehat{\text{Var}}[I(k)] - \text{Var}[I(k)]| \leq \sum_{i=0}^k (i - E[I(k)])^2 |\Delta p_i| + (\Delta E[I(k)])^2,$$

by substituting Δp_i with $|\Delta p_i| \leq \epsilon \sqrt{\Pr[I(k) = i]}$ we have

$$|\widehat{\text{Var}}[I(k)] - \text{Var}[I(k)]| \leq \sum_{i=0}^k (i - E[I(k)])^2 \epsilon \sqrt{\Pr[I(k) = i]} + (\Delta E[I(k)])^2,$$

by applying Cauchy's Inequality we have,

$$\begin{aligned} |\widehat{\text{Var}}[I(k)] - \text{Var}[I(k)]| &\leq \epsilon \left(\sum_{i=0}^k (i - E[I(k)])^2 \Pr[I(k) = i] \right)^{1/2} \left(\sum_{i=0}^k (i - E[I(k)])^2 \right)^{1/2} + (\Delta E[I(k)])^2 \\ &\leq \epsilon \left((k+1) \sqrt{(2k+1)\text{Var}[I(k)]/6} \right) + (\Delta E[I(k)])^2, \end{aligned}$$

the proof can be completed by substituting $(\Delta E[X])^2$ with Inequality (36). ■